

The Civic Epistemology of Liberal Democracy

Dissertation

Presented in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy
in the Graduate School of The Ohio State University

By

Jonathan David Whitsett, B.A.

Department of Political Science

The Ohio State University

2021

Dissertation Committee:

Michael Neblo, Adviser

Eric MacGilvray, Co-Adviser

Tom Wood

Copyright by
Jonathan David Whitsett
2021

Abstract

The internet and social media have profoundly altered the informational environment of liberal democracies in the 21st century. Much of the response to these developments by political science scholars and mainstream political commentators is pessimistic and reflects a deep pessimism about the judgment of lay citizens that is at odds with support for liberal democracy. This pessimism, in turn, has driven many to embrace illiberal public policies and views of good citizenship. I argue that these trends are both unjustified and worrisome, as they are likely to undermine society's collective ability to improve its knowledge and discover illegitimate authority. Instead of viewing the judgment of lay citizens as a liability to liberal democracy, we ought to view it as one of its foundational assets. Accordingly, we ought to view the impact of 21st century communications technologies in a more optimistic light and support public policies and models of good citizenship that encourage citizens to engage in the informational environment with assertiveness and confidence.

Dedication

To Jack, who helped me see the value of gadflies

Acknowledgments

Like many students before and no doubt many to come, I owe a debt of gratitude to my dissertation advisers, Michael, Eric and Tom, all of whom have consistently challenged me and taught me to demand a higher standard of justification for my beliefs. I am also grateful to have landed in the political science department at Ohio State, which has consistently given me a combination of support and freedom that is every graduate student's dream. I hope I have managed to at least partly justify that support and freedom with this dissertation.

Vita

May 2003.....West Liberty-Salem HS
2010.....B. A. Political Science, University of Minnesota-Morris
2015 to present.....Graduate Associate, Department of Political Science,
The Ohio State University

Fields of Study

Major Field: Political Science

Table of Contents

Abstract	i
Dedication	ii
Acknowledgments	iii
Vita.....	iv
List of Figures.....	vi
Introduction: Progressivism, Pessimism, and the Post-Truth “Malaise”	1
Chapter 1: A Democracy Worth Defending	34
Chapter 2: The Epistemic Well-Being of Democracy	57
Chapter 3: Individual Vice, Collective Virtue	119
Chapter 4: Civic Epistemology	176
Chapter 5: Consilience, Confidence, and Justified Deference	228
Chapter 6: Epistemic Pluralism and 21 st Century Citizenship	304
References	367

List of Figures

Figure 1. Vaccine Coverage in the United States, 2012-2016.....	142
Figure 2. Vaccine Coverage in the United States, 2007-2018.....	142

Introduction: Progressivism, Pessimism, and the Post-Truth “Malaise”

“Faith in democracy and in the market was shaken. With the questioning of...central institutions, there was also questioning of the implicit vision that independent individuals voting in a democracy would produce the right decision, and that independent individuals uncovering facts in a random fashion would reveal truth...That, at least, was the early perception, the view of the Progressive Era.”

-Michael Schudson¹

Section One: The Legacy of Progressivism

Progressivism in the Zeitgeist

In American² history classes typical of those taken by high schoolers and university undergraduates, Progressivism is little more than a footnote. If it is addressed at all, it is more as a means of describing the end of the era it replaced – that is, that hazy era lying between the end of the Civil War and the Great Depression, which is viewed by many as a rather embarrassing moment in American history rife with corruption, graft, corporatism and yellow journalism – than the one it ushered in. For those scattered history buffs and abnormally well-read citizens who go beyond the contents of their high school history courses and look into the era further, the Progressive age is generally viewed in a positive light, having (for example) replaced party machines and pre-printed “slip tickets” with anti-corruption laws and the Australian ballot and replaced the scandalous sensationalism of yellow journalism with the respectable, professionalized version of journalism taught in many American universities and synonymous, for most of us, with the term “the news.” For the well-read few who know much of

¹ Schudson (1978, 158)

² Here and throughout the dissertation, I focus on the American political context. I do so because that is the context with which I am by far most familiar. However, I believe that in our globalized era many of the issues with which I grapple in this dissertation apply to many other polities throughout the world, and the theoretical arguments I pose apply to any context in which the principles of what I call liberal democracy (defined below) are at stake.

Progressivism, these are the essentials that sum up its legacy. Most, however, could not even identify these few essentials as hallmark features of Progressivism.

Yet, as ignorant as most of us are about the Progressive era, many of us in the contemporary era, including many outside the borders of the United States itself, are very much the heirs of Progressivism when it comes to our views of our expectations of democracy and the type of citizenship appropriate to it. Many of us believe, for example, that people's vote ought to be cast on the basis of public policy considerations rather than fanatical allegiance to whatever political party dominates our family or our geographic community. Many of us also believe that it is a civic duty to remain informed about politics, and that the best way to remain informed in this way is to stay abreast of the news. And many of us believe that not just *any* news will do for the purposes of fulfilling this duty, but instead that only *some* news is of a high enough quality to meet the demands of democracy – namely, news produced by reporters and editors who have received specialized training and are employed by a specific subset of professionalized news producing companies. For many of us, indeed, only news that meets such a standard is “fit to print.” In these and many other respects, our default views about democracy, knowledge, and democratic citizenship are a direct inheritance from, or at least a reflection of, Progressive views on the same.³

Disillusionment about Democracy as a Legacy of Progressivism

But we are also the heirs of Progressivism in another way, one that can be seen in the epigraph printed at the beginning of this essay, which comes from historian Michael Schudson. For the legacy of Progressivism is not just a legacy of vanquishing corrupt party bosses and scurrilous yellow journalists from the scene of American politics. At least, that was not the way it was seen by those who followed immediately in the Progressive Era's wake. On the contrary, for many of their immediate successors, the legacy of Progressivism was a legacy of *disillusionment about democracy*, especially in terms of the performance of their civic duties by lay citizens. The aims of Progressives were high. They wanted a highly informed citizenry remaining on top of public affairs and eschewing partisanship and self-interest to vote in the public's interest on the

³ For an excellent account of how Progressivism has impacted our views about citizenship, see Schudson (1998, ch. 4).

basis of public policy merit. The results of the types of policies Progressives headed up, however, have generally been interpreted as mediocre at best. To be sure, Progressive initiatives effectively reduced graft, did away with the pre-printed “slip ticket” and ushered in the decline of yellow journalism, but at the same time participation in popular elections plummeted from all-time highs of over 80% to the now-familiar levels of just over half the eligible voting public. Citizens who embraced the view of citizenship favored by Progressives found themselves daunted by the immense moral and logistical expectations placed on them by that view. Many retreated from politics and still more failed miserably to live up to the Progressive standard. Post-Progressive critic Walter Lippmann expressed sympathy for both failures to live up to that standard in a passage of *The Phantom Public* that is worth repeating in full:

“There is then nothing particularly new in the disenchantment which the private citizen expresses by not voting at all, by voting only for the head of the ticket, by staying away from the primaries, by not reading speeches and documents, by the whole list of sins of omission for which he is denounced. I shall not denounce him further. My sympathies are with him, for I believe he has been saddled with an impossible task and that he is asked to practice and unattainable ideal. I find it so myself for, although public business is my main interest and I give most of my time to watching it, I cannot find time to do what is expected of me in the theory of democracy; that is, to know what is going on and to have an opinion worth expressing on every question which confronts a self-governing community. And I have not happened to meet anybody, from a President of the United States to a professor of political science, who came anywhere near to embodying the accepted ideal of the sovereign and omniscient citizen.”⁴

Citizens' realization of the daunting demands implied by the Progressive ideal of citizenship, along with democratic citizens' routine failure to live up to it, meant that for many, like Lippmann, who lived through the era during which Progressives' most famous reforms were put into effect, the legacy of Progressivism was disillusionment about democracy. In part, this is

4 Lippmann (2012 [1925]), 20-21

because the fallout from the Progressive movement was interpreted by many commentators as justifying a deep degree of pessimism about lay citizens, particularly their capacity to render reliable judgments about truth both in general and in the political arena. As Schudson writes, in the post-Progressive era, “[f]aith in democracy was losing out to fears of unreasoning – and of the presumably unreasoning...masses.”⁵ Political observers and intellectuals of the time expressed a profound “distrust, not so much of reason as of the public's capacity for exercising it.”⁶ Lippmann cited the failure to achieve the Progressive ideal of citizenship as having “produced the current disenchantment”⁷ with democracy (though he himself advocated radically revising our views of democracy rather than doing away with our commitment to democracy itself).⁸

Institutionalization and Authority as Responses to Democratic Disillusionment

As history would have it, the post-Progressive response to their disillusionment about lay judgment – and by extension democracy – was not to *abandon* democracy, a fact for which we may have the “democratic West against the authoritarian East” moral framework through which World War II and the Cold War were interpreted by the West to thank, but instead to *institutionalize* it. Luminaries of the day like Lippmann and his counterpart Edward Bernays advocated the development of institutions designed to help limit the damage that could be done to democracy by the public's deficiencies in judgment in order to put into practice better public policies than should be expected to garner public support through the chaos of an unmanaged “marketplace of ideas.” This was the public-minded justification Bernays gave for his work developing the industry of public relations,⁹ while for his part Lippmann advocated for the development of institutions to serve this purpose both in the form of government agencies for managing information¹⁰ and professional standards for the practice of journalism.¹¹ These moves by Lippmann and Bernays were typical of the West's response to its post-Progressive disillusionment with democracy, which saw intellectuals and societal influencers from World

5 Schudson (1978), 131

6 Ibid., 129

7 Lippmann (2012 [1925]), 39

8 Ibid., chs. 4 & 5

9 Bernays (2005 [1928]), ch. 6

10 Lippmann (2012 [1925]), Part VIII

11 Lippmann (1920)

War II all the way through the Cold War regularly advocate for expansions in the roles of institutionalization, professionalism and expertise both in governance through the increasing role of scientists, economists and other experts, and in the gathering and distribution of information through an increasingly professionalized news media. By increasingly portraying these figures and institutions as integral to the well-being of democracy while simultaneously holding deep doubts about the fitness of lay citizens to epistemically navigate an increasingly large, complex and interconnected society, these luminaries came to embrace an image of responsible democratic citizenship that largely consisted of lay citizens adopting an attitude of routine deference to authorities bearing the trappings of belonging to the right kinds of institutions, having received the right sorts of professional training, and/or bearing the appropriate forms of credentials to demonstrate their expertise, and declining to invoke their own judgment as grounds for disagreeing with the conclusions of these authorities.

In this sense, too, I believe many of us – especially academics in the social sciences, journalists, and many others who would be described as “elites” in the political science jargon – are very much the heirs of Progressivism, or at least the heirs to the legacy of Progressivism. That is to say, our default views about responsible citizenship often include citizens being willing to routinely trust the claims of scientists who belong to institutions like NASA, the reports published by news agencies that employ professional journalists over those that do not, and the information passed on by other established institutions that have come, through some process that is seldom if ever articulated to lay citizens themselves, to be dubbed “reliable sources.” This legacy of Progressivism could be seen in grade school curricula of those of us who grew up in in the pre-internet era and in the early days of the internet, where we were taught that legacy newspapers, books found in certain sections of the school library, and things published in academic journals counted as “reliable” sources of information while, while sources that did not bear the same marks of institutional accreditation were not. In many cases (including my own), this was the entirety of our training in what I in this dissertation will call *civic epistemology*, that is, the process of forming beliefs and sharing information appropriate to democratic citizens. Many students today, including many of those who have stumbled into my contemporary issues in American politics class, are taught a similar approach to the one I was taught several decades ago, updated slightly to apply to newly emergent 21st century communications technologies.

Students today are taught that legacy news media are more reliable than internet-era startups like *Vice* and *Buzzfeed*, and *certainly* more reliable than most contents found on *Wikipedia* and *certainly* more reliable than practically anything shared on social media. Those few students who report being taught anything beyond this are usually taught *both* that the above-mentioned sources are the *most* reliable ones and, should they foray any farther afield in their research than these, to look for signs of institutional officialdom, such as the existence of a “.org” or “.gov” or “.edu” suffix to the site's URL, placing their trust in those sources that bear such institutional markers and mistrusting those that do not. Beyond the classroom, too, this preference for the same specific, easily identifiable, long-established institutions like those of the scientific community and the legacy news media dominates our thinking about which sources ought to be considered “reliable.” In my experience, this is especially true of elites like those mentioned at the beginning of this paragraph.

This centrality to our default notions of civic epistemology of citizens being willing to trust in sources whose reliability is defined by their belonging to a certain set of established institutions is a legacy of Progressivism. In this dissertation, I hope to clearly demonstrate that, like that legacy, it is subsidized by a *deep pessimism about the judgment of lay citizens*. This deep pessimism was often explicit for the original heirs of Progressivism, but for us it is more often implicit. Yet, it is pervasive. It is implicit in the scorn and suspicion with which we view user-driven information-sharing platforms, like social media and Wikipedia. It can be seen in our ubiquitous fears about the vulnerability of democratic citizens to manipulation, misinformation and demagoguery. It lies behind many commentators' encouragement of citizens to practice good citizenship by seeking out the claims of established institutional authorities and accept what *those authorities say*, rather than simply instructing them to seek out a *wide breadth* of information and *judge for themselves* what to believe. And it is implicit in the regular and vocal condemnations that spring forth from the pages of academic journals and editorials whenever citizens choose, on the basis of their own judgment, to reject the claims of these established institutional authorities. Just as we have inherited a legacy of pessimism about lay citizens' judgment from Progressives, we have inherited the solution that won out following the Progressive era. That is, we have sought to safeguard democracy from the suspect judgment of

the masses by turning to the judgment of institution. For many, it is on *these* institutions, and the public's willingness to place its trust in them, that the well-being of democracy truly depends.

The Heirs of Progressivism and the Post-Truth “Malaise”

A final sense in which many of us are heirs of Progressivism is that, just like the Progressives, we live in an era where revolutionary developments that at first raised optimism about the prospects of democracy – for Progressives, this was an industrial revolution that made widespread material prosperity and control by the masses of their own collective destiny seem more within grasp than ever; for both Progressives and their heirs, it was the increasing attainment of unprecedented levels of education; for us, it is an information revolution that seemed to promise a dramatic expansion of the power of the voices and choices of the masses – are now being widely interpreted as grounds for deep pessimism about the judgment of lay citizens and by extension disillusionment with democracy.

The depth of this pessimism at the current moment is very widespread and often profound. Indeed, it is so common, and so commonly the topic of discussion, that commentators have found it convenient to develop a shorthand to summarize what they take to be the depraved epistemic state to which the public has fallen. This is the term “post-truth,” whose many uses are well illustrated by Stephan Lewandowsky and colleagues in several recent articles found in a special “post-truth” edition of the *Journal of Applied Research in Cognitive Memory and Development*. In those articles, Lewandowsky and colleagues characterize the post-truth condition as a “malaise,” a “crisis” and a “problem.”¹² In their view, the post-truth “malaise” entails some combination of the rejection of established expertise in favor of the “rule of the well-financed or the prejudices of the uninformed;” the “prevalence of misinformation,” which they view as dominating the online environment;¹³ the descent of citizens into “alternative epistemologies” that are “not easily punctured by empirical evidence or corrections issued by ‘elitist’ media or politicians;¹⁴ and possibly a growing tendency for some citizens to “stop believing in facts altogether.”¹⁵ The sentiments of these scholars are mirrored in the mainstream

¹² Lewandowsky et al (2017a, 2017b)

¹³ Lewandowsky et al (2017a), 354

¹⁴ Ibid, 356

¹⁵ Ibid, 355

political commentary sphere. Typical is the opinion voiced by Julian Birkinshaw in a recent *Fortune* article, who proclaimed that we “are living in a post-truth world,” which for Birkinshaw means that “alternative facts and fake news compete on an equal footing with peer-reviewed research and formerly-authoritative sources such as...the BBC.”¹⁶ As Johan Farkas and Jannick Schou note, this narrative of the post-truth condition as having thrown democracy into “a state of crisis, with fake news flooding the Western world and alternative facts breaking down the very core of decision-making” has become “one of the dominant portrayals of present-day democratic societies.”¹⁷

Just like the post-Progressive era, then, our era is one in which pessimism about the judgment of the public is widespread. Oddly enough, at the very same time this pessimism toward the *public* is often expressed quite explicitly in our era as in that of post-Progressives, we live in an era where, unlike the post-Progressive era, it is impolite – even scandalous – to explicitly disavow *democracy*.¹⁸ It is impossible to tell whether it is good manners that keeps the same commentators who express such deep pessimism about the judgment of lay citizens from voicing doubts about the viability of democracy but, raise them or not, it is equally impossible to deny that such doubts must necessarily arise from such a pessimistic view of the public. Though they are nearly never expressed publicly as such, I believe grave inhibitions about democracy are in fact very widespread. Curiously, they are nearly always phrased as anguished pleas on democracy's *behalf*. For instance, a large number of scholars have voiced alarm at the rise of “illiberal democracies” in which the public's prejudices and susceptibility to demagogues is seen as steering democracy to embrace authoritarian leaders and repressive public policies.¹⁹ The pervasiveness of this tendency to connect the cognitive and moral deficiencies of the public to demagoguery and bad political outcomes is also evident in the near-ubiquitous tendency to use the term “populism” as a synonym for “authoritarianism.” Indeed, legal scholar Richard Posner points out that while there is no shortage of defenses for elitist models of democracy, populist democracy is a position on the “democratic spectrum that has almost no support among political

16 Birkinshaw (2017)

17 Farkas and Schou (2020), 45

18 Scandalous though it may be, concerns like those identified here have driven some theorists, such as Jason Brennan, to disavow democracy explicitly regardless of the opprobrium that move is likely to incur from their peers (Brennan 2016).

19 See, e.g., Zakaria (1997), Wodak (2019)

theorists.”²⁰ It is no stretch to say that in our current populism-averse era the applicability of this statement extends outside the musty confines of the offices of political theorists. Closer to the truth would be to say that many scholars and commentators agree with the paradoxical statement made by Stephen Levitsky and Daniel Ziblatt in their iconically titled *How Democracies Die* that “Democratic backsliding today begins in the ballot box.”²¹

That so many commentators see the ballot box itself as one of the primary *threats* to democracy is indicative of the curious mix of pessimism about the public and support for democracy that reigns in the current moment. Yet, like Levitsky, Ziblatt, and these others, most of today's heirs of Progressivism combine grave doubts about the reliability of the public's judgment with the wish for democracy to succeed. By expressing their continued (if wishful) support for democracy, they seem to differ from the post-Progressive era mentioned earlier, where the likes of Carl Schmidt were beginning to compose direct attacks on the viability of democracy and even many democratic theorists, such as Lippmann and followers of Joseph Schumpeter, only managed to maintain their support for democracy by insisting on a revision of our expectations of it so radical that many of those who read them today find them scarcely to be defenses at all.²² But given the depth of the pessimism they express about the public, it is not easy to see why contemporary commentators' attachment to democracy should be so strong. Indeed, I will argue in Chapter One that this level of pessimism and full-throated support of democracy – at least, any sort of democracy *worth* supporting – are fundamentally incompatible, and that one must do away with either the pessimism or the democracy. In this dissertation, I will advocate for the former, but this is not the usual move made by the heirs of Progressivism today. Instead, they tend to make the same move Western societies made in the wake of their disillusionment with Progressivism, that is, by seeking to put in place some set of *authorities*

20 Posner (2003), 155

21 Levitsky and Ziblatt (2018), 5

22 Lippmann's *The Phantom Public* is usually read more as an *attack* on democracy than a normative *defense* of it (though I think that is a mistake), while most interpret Schumpeter's *Capitalism, Socialism, and Democracy* – this time, in my opinion, correctly – as a *descriptive* account of democracy that is not even explicitly normative at all. However, Schumpeter has been used by political theorists as an account of a kind of democracy worth defending. Posner, for instance, characterizes his account of “Concept 2 democracy” in *Law, Pragmatism, and Democracy* as “recasting [Schumpeter's account] as a normative rather than a merely positive concept” (Posner 2003, 165).

whose task it is to *manage* democracy from above, and cultivating in the public an *attitude of rote deference to those authorities*.

Information Managerialism and Routine Deference to Authority as Progressive Responses to Pessimism about the Public

Walter Lippmann's *Public Opinion*, first published exactly a century ago as of this writing, is a masterpiece that remains, to my knowledge, the single best examination of the epistemic plight faced by contemporary democracies ever written. Lippmann's description of the circumstances typical of democratic societies under conditions of modernity – the overwhelming size and complexity of modern life, the profound limitations of individuals' epistemic tools for evaluating public policies and public figures, the at times counterintuitive dynamics of *collective* versus *individual* knowledge collection, and countless other issues – is masterful, and many of the conclusions Lippmann draws about the problems of relying on a politics driven by “public opinion” under such conditions are unavoidable. As such, Lippmann will be cited at some point in every chapter of this dissertation. Given the perspicacity of Lippmann's account in these regards, it should be no surprise that public policy recommendations similar to those that constitute the final few chapters of *Public Opinion* often appeal also to other observers of democratic affairs who become aware of some of the problems Lippmann so excellently identified a century ago. And indeed the writings of those most worried about the post-truth “malaise” imply agreement with Lippmann on multiple fronts. Specifically, they respond to the difficult epistemic conditions faced by democratic citizens under conditions of modernity by (a) restricting the range of affairs on which it is deemed appropriate for lay citizens to confidently rely on their *own judgment* to form appropriate conclusions and (b) advocate that in all other areas but this restricted range the appropriate thing for citizens to do is to *defer to established authorities*.

Tendencies (a) and (b) are ubiquitous in the literature warning of the threat posed to democracy by post-truth, where one can see accusations that take the following basic structure hurled out again and again in a seemingly endless stream of permutations. Some subset of the public will be identified who insist on coming to their own conclusions about what is true or which authorities ought to be trusted in a given epistemic domain, such as climate science or

immunology. The commentator citing this example will name what they take to be the *clear* authorities on the matter and express incredulity at the very idea that such a subset of the public could ever reasonably refuse to defer to those authorities. The commentator takes such behavior to be not only self-evidently irrational, but in fact *so* irrational as to be a sign of some kind of widespread epistemic malignancy spreading through the body politic.²³ One of the epistemic authorities to which deference is most frequently expected, and its denial most frequently condemned, is the vaguely defined “scientific community.”²⁴ Another is the legacy news media. The refusal to defer to either of these sets of established epistemic authorities by some subset of the population is cited by numerous scholars as evidence of democracy-threatening levels of epistemic vice among the public.²⁵ In the overwhelming majority of cases in which the epistemic authority of the scientific community is invoked, including all those cited above, no effort whatsoever is made to present an *argument why* citizens ought to defer to those particular authorities on the matter at hand even when the citizens' own judgment suggests they either come to a different conclusion or place their trust in some set of authorities not favored by the commentator. Instead, the commentator simply expects *that* their co-citizens ought to defer to such authorities as a matter of course. In other words, the commentators expect not *just* deference but *routine deference* to their preferred authorities. The regular adoption of such an attitude toward those authorities is viewed by them as a sign of *good citizenship* and refusal to do so a *threat to democracy*.

The same tendency can be seen among these commentators in the disposition they advocate toward the legacy news media vis-a-vis informational sources that have sprung up in the internet era. Almost unanimously, the commentators who express the greatest fears about the post-truth condition express the greatest faith in legacy journalism. Anya Schiffrin, for example, ranks the legacy news media alongside institutions of “education” and “science” – and, notably,

23 Indeed, the metaphor of disease is prevalent in this discourse, with some articles about the post-truth “condition” even being published in public health journals, (e.g. Speed and Mannion 2017). For a detailed account of just how prevalent this tendency is, see Farkas and Schou (2020), 46-49.

24 I put the term in scare quotes here because I think it encompasses many very different intellectual traditions, using very different techniques, whose average levels of reliability in terms of the knowledge they produce and disseminate vary greatly. Because of this wide variability, I object to the way members of this “community” are often treated as equivalents in terms of the readiness with which lay citizens should be prepared to defer to their claims. I'll continue to use the term throughout the dissertation without the scare quotes, but I do so chiefly because Progressives often use the term themselves in this way, and not because I endorse that usage.

25 E.g. Lewandowsky et al (2017), Schiffrin (2017), Deb, Donohue and Glaisyer (2017)

not the discretion of the public – as three of the main aspects of democratic societies that “traditionally served to keep ‘false facts’ and demagoguery at bay.”²⁶ The view that established news companies are both more reliable than online-era alternatives and integral to maintaining the health of democracy is as pervasive among those who see themselves as the enemies of post-truth as the view that the scientific community deserves the routine deference of citizens in the epistemic domains commentators see as belonging to that community.²⁷ So is commentators’ proclivity for assertion over argument. When most commentators write about the democratic value of “freedom of the press,” for example, it is clear that they mean “the press” to be synonymous with “professional journalists,” and *not* online message boards like Reddit, user-generated content like that typical of Youtube, or people’s Facebook or Twitter feeds. It is only the *former* – those who belong to an established institutional authority – they view as integral to the health of democracy; the latter is a *threat* to it. Commentators worried about post-truth often simultaneously cite declining rates of trust in legacy news media as a bad sign for democracy and condemn citizens who turn instead to online alternatives. Once again, implicit in such worry and condemnation is the idea that democracy would be better off if citizens simply trusted the claims of some set of established authorities in a given epistemic domain – this time the news – rather than using their own judgment to decide what’s true upon sampling the much wider range of information alternatives available on the internet.

Commentators who view 21st century communications technologies as a threat to democracy often lump these two sets of authorities – the scientific community and the legacy news media – together as paradigmatic examples of the sorts of institutions to whose claims citizens ought to be willing to routinely defer in order to preserve the well-being of democracy. In this dissertation, I will argue that this expectation of *routine deference* to authorities whose legitimacy is defined prior to and outside of any act of judgment by the lay citizen makes perfect sense given commentators’ pessimism about the judgment *of* lay citizens. Choosing the *right* authorities to defer to in the *right* situations would appear to necessarily entail a sound exercise

26 Schiffrin (2017), 123

27 Very often, in fact, support for the legacy news media among opponents of post-truth is taken for granted to the point that it is not even explicitly asserted at all. Instead, the commentator simply takes for granted that we all *already* agree that legacy news media are more trustworthy than online alternatives. For instance, the British House of Commons’ Interim Report on Misinformation and “Fake News” uses the term “traditional” as synonymous with “reliable” in contrasting “traditional forms of communication” with social media, while clearly depicting the latter as a vector for “false, misleading, and persuasive content” (House Report 2018, 4).

of judgment, but that is precisely the faculty about which those worried about the post-truth condition have the gravest doubts. The only practicable alternative would seem to be to remove important decisions about whose claims to believe on matters of central importance from the untrustworthy public and place them somewhere else more reliable, then train the public to “know its place,” only using its judgment in that limited domain it can understand and leaving the rest to the authorities. This is a plausible conclusion from the work of Lippmann,²⁸ who in *Public Opinion* emphasized how constrained the range of affairs over which the public was capable of issuing a reliable judgment was and in *The Phantom Public* argued in no uncertain terms that on most political issues the public's judgment is fit to do little except “meddle ignorantly or tyrannically,”²⁹ and whose general approach to the epistemic and other dilemmas he saw facing modern democracies was a combination of designing sound institutions and placing decision-making in the hands of elites. Thus, many of those who are most worried about the post-truth movement may be considered heirs of Progressivism in the sense that, like Progressivism's more immediate heirs such as Lippmann, their pessimism about lay citizens' judgment drives them to attempt to envision a form of democracy that depends very little on it – in fact, depends largely on teaching citizens *not* to exercise it, and instead encourages them to defer to those whose epistemic reliability is presumed to be superior to that of citizens themselves.

Progressives' Self-Identification as Defenders of Enlightenment Ideals, Liberalism and Democracy

In addition to defenders of democracy, opponents of post-truth often also view themselves as proponents of Enlightenment ideals of reason, evidence, and belief, accusing post-truthers as opposing all of these things. Communications scholar Peter Dahlgren, for instance, portrays the post-truth movement as consisting in a series of “aggressive attacks on basic

28 I am not really sure whether it was Lippmann's own view. In *Liberty and the News*, Lippmann's main emphasis is on making news organizations adhere to quasi-scientific standards and rely only on “objective facts,” but toward the end he also mentions the need for “drastic competition [from] those whose interests are not represented in the existing news-organization” (Lippmann 1920, 60). In later works, however, particularly *Public Opinion* and *The Phantom Public*, Lippmann's emphasis is increasingly on the deficiencies of public judgment and the need to move toward more institutional bases for public decision-making. I will turn to this discussion in more depth in the final chapter of this dissertation.

29 Lippmann (1925), 70

Enlightenment premises” and its members as subscribing to a “new epistemic regime, where emotional response prevails over factual evidence and reasoned analysis.”³⁰ The last part of this accusation is reminiscent of the *Oxford English Dictionary's* oft-cited definition of post-truth as “circumstances in which objective facts are less influential in shaping public opinion than appeals to emotion and personal belief.”³¹ Elsewhere, psychologist Stephan Lewandowsky and colleagues characterize post-truth epistemology as not conforming to “conventional standards of evidentiary support,” “not easily punctured by empirical evidence,”³² and encouraging citizens to “stop believing in the existence of facts altogether.”³³ Jason Hannan agrees that post-truthers subscribe to a “curious epistemology” where “[p]opularity now competes with logic and evidence as an arbiter of truth,” and joins many anti-post-truthers in identifying social media as a hotbed of and contributor to that epistemology, which he clearly opposes.³⁴ In the mainstream commentary sphere, the editorial board of *The Economist* accuses post-truthers of subscribing to the idea “that truth is not falsified, or contested, but of secondary importance.”³⁵ These examples illustrate how post-truth opponents in both the academic and mainstream worlds see themselves as defenders of Enlightenment ideals regarding reason, evidence, and belief, while viewing the post-truth movement as presenting an alternative epistemology in opposition to those ideals.

Anti-post-truthers also often see themselves as proponents of *both* liberalism and democracy. Their support for democracy is sometimes explicit but more often it is implicit in the way, as pointed out above, they tend to frame their arguments as warnings of the threat posed by the post-truth movement and the evolution of the 21st century informational environment to democracy, which presumably would not be worrisome if democracy were not something worth defending. Their support for liberalism, meanwhile, is implicit in their defense of what I will call core liberal *freedoms of communication*, such as freedom of speech and freedom of the press, and view both the post-truth mentality and emergent communications technologies, especially social media, as posing a threat to them. Schiffrin, for instance, seems to embrace free speech when she advocates against letting policies intended to fight fake news online “become an excuse for

30 Dahlgren (2018), 25

31 *BBC News* (November 11th, 2016)

32 Lewandowsky et al (2017), 356

33 *Ibid.*, 355

34 Hannan (2017), 220

35 *The Economist* (September 9th, 2016)

corporate and government censorship,” while in the same breath warning of the danger to democracy of allowing “technology companies use free speech as an excuse not to take action.”³⁶ Elsewhere, in an article arguing that lay citizens should consider the legacy news media to be a more reliable source of information than the “blogosphere” and that the growing tendency for people to distrust the legacy news media in favor of alternative sources online is bad for democracy, epistemologist Alvin Goldman endorses the common truism that “democracy requires a free press” because “only a free press can ferret out crucial political truths and communicate them to the public.”³⁷ Similarly, as we will see, critics of phenomena frequently attributed to the post-truth movement by its opponents, such as “fake news,” often cite a desire to see democracy flourish or to preserve liberal values.³⁸

Civic Epistemology

Up to now, I have used the term “Progressive” to refer to the American political movement going by that name at the turn of the 20th century. I have briefly traced the intellectual legacy of Progressivism through its immediate heirs to the present and identified that intellectual legacy as one that includes a *deep pessimism* about the judgment of lay citizens and encourages citizens to adopt an attitude of *routine deference to established intellectual authorities* as an important facet of citizenship in modern democracy. From now on, I will use the term “Progressive” to describe not the original Progressives but the contemporary opponents of the post-truth movement who have inherited that intellectual legacy. Though that legacy is certainly not the one most Progressives thought they would be handing down to their descendants, enough of Progressivism in its original sense – for instance, its heavy emphasis on the importance of professionals and experts, its views of the proper form and function of the news media, and its ideal of “informedness” as integral to sound citizenship – lives on in the work of contemporary opponents of post-truth to justify using the same word to describe these contemporary commentators. From here on out, then, I will use the term Progressive to refer to this array of contemporary opponents of the post-truth movement.

36 Schiffrin (2017), 122

37 Goldman (2010), 112

38 E.g. *Boston Globe* (2017)

Drawing together the aspects of the Progressive position described above, we can begin to piece together a *Progressive Model of Civic Epistemology*. The term *civic epistemology* is not yet a commonly used term – in fact, I am not aware of its being used on a regular basis by any political commentator to date – but I believe it to capture a vital concept endemic to much contemporary political discourse that is importantly different from the more familiar terms *social epistemology* and *epistemic democracy*. The main foci of these latter two terms are to *examine the dynamics of collective knowledge acquisition*, which I take to be the main priority of social epistemology, and to *explore whether, to what extent, and in which forms democracy is justified on an epistemic basis*, which I take to be the central topic of epistemic democracy. While *civic epistemology* is concerned with such collective epistemic outcomes in the sense that one of its goals is to help encourage citizens to behave in ways that foster beneficial epistemic outcomes, its central concern is not merely to *describe* those outcomes or to *defend democracy* as a political system. Instead, it is to *develop normative models of belief formation and communication for democratic citizens*. In my view, this is best done by trying to simultaneously hold three things in mind – (1) the epistemic circumstances faced by lay citizens trying to form reliable beliefs; (2) the likely outcome(s) of the epistemic behaviors practiced by individual citizens at the collective level; and (3) core democratic values, including but not limited to maximizing knowledge acquisition – and developing models of citizenship that strike the best balance between maintaining fidelity to those core values and helping to maximize collective knowledge acquisition while keeping in mind the on-the-ground realities faced by democratic citizens.

Civic epistemology is clearly an implicit concept in many current-day political discussions, including those emphasized by Progressives in the examples given so far. Indeed, it is a *central* concept in a great many of those discussions. It is implicit in debates about the confidence we should place in climate scientists' claims about the causes and likely effects of climate change. It is implicit in debates about the ethicality of refusing to vaccinate one's children. It is implicit in debates about what, if anything, Facebook, Twitter and other social media platforms should do about fake news. It has played a central role in the discussion over the United Kingdom's referendum sanctioning their exit from the European Union, or “Brexit.” And it is implicit in the myriad debates about the proper role and scope of expertise in democratic society, as well as the even more numerous debates about the epistemic effects of the internet on

democracy and the public policies we ought to take in response to them. All of these topics of political discussion rely on, and have implications for, models of the epistemic and communicative behaviors appropriate to democratic citizens in the 21st century. Given its centrality in all these discussions, all of which have significant implications for the ways we approach public policy and, consequently, the shape taken by democracy in the future, it is important that we make sure our notions of civic epistemology are built on a solid foundation. I believe Progressives have thus far failed to do so. Because of this, they have regularly recommended public policies and the adoption of epistemic habits and dispositions on the part of citizens that are destructive both to the collective pursuit of knowledge and to the vital societal task of maintaining vigilance toward authorities. Articulating the views implicit in the Progressive position, fleshing out the model of civic epistemology implicit in them, identifying its problems, and providing a better model are the main aims of this dissertation.

Part Two: Civic Epistemology in the 21st Century

The Civic Epistemology of Progressivism: Deferential Democracy

I believe it is obvious that Progressives of the sort introduced in this chapter share an implicit model of civic epistemology. The key aspect of that model I want to emphasize is its *veneration of* and advocacy of *routine deference to (certain of Progressives' favored) established epistemic authorities*. The epistemic authorities to which routine deference is most commonly endorsed by Progressives, and therefore the two that will serve as my focus in this dissertation, are the scientific community and the legacy news media. Progressives' veneration of these two established epistemic authorities is evident in the near-universality with which they defend their reliability and disparage citizens who place their trust anywhere else. The expectation of *routine deference* to their claims is implicit in the propensity for Progressives to treat *the very fact of* refusal to defer to these established authorities' claims as *prima facie* evidence of unreason, along with the consistent absence of any effort whatsoever by those same commentators to *explain why* that expectation is justified. In none of the examples above, which are typical of Progressive disparagements of the post-truth movement, is any effort made by the author(s) to supply any sort of reasoning to *justify* the demand of deference to the named authorities. Instead, in all cases

the authors simply *assert that* such deference is appropriate and beneficial to democracy and that, *ipso facto*, refusal to do so is inappropriate for lay citizens and detrimental to democracy. As I will explain in greater detail in Chapter One, both liberalism and democracy heavily emphasize the need for demands of deference to authority to be *justified* to those citizens whose deference is demanded. The implication of Progressives' consistent tendency to demand citizens' deference without presenting any such justification would seem to be that there ought to be a *default* presumption in favor of deferring to the claims disseminated by the scientific community and the legacy news media (among other of Progressives' favored established epistemic authorities), a presumption so strong that refusals to defer are either *definitionally unreasonable* or, at the very least, bear the burden of proof, so that refusals to defer must be justified in order to be consistent with responsible epistemic citizenship. This is why I characterize Progressives as not just advocating deference but *routine* deference to (favored) established epistemic authorities, such as the legacy news media and the scientific community. Because the expectation of routine deference to these authorities is so central to Progressives' criticisms of post-truthers' and warnings of democratic decline, I will sometimes call their model of civic epistemology *deferential democracy* for short.

For reasons I have already hinted at, I will presently argue that there is considerable tension between Progressives' endorsement of *deferential democracy* and their avowed support for Enlightenment ideals and liberal democracy. However, it is vital to first recognize that their recommendation of routine deference to established authorities and institutions is consistent with at least two other aspects of the Progressive point of view. It is, first of all, as I argued above, consistent with the *deep pessimism* they regularly espouse about the judgment of lay citizens. Indeed, I believe Progressives' view that good citizenship consists largely in routinely deferring to the claims that disseminate from established epistemic authorities stems largely *from* that pessimism. Progressives view the current informational environment characterized by high choice and low barriers to publication as currently leading to, and bound to result in, a state of affairs where fake news and misinformation compete on an equal footing with truth,³⁹ demagoguery and manipulation reign,⁴⁰ and the public bases its beliefs not on evidence and

39 Silverman (2016), Vosoughi et al (2017), Lewandowsky et al (2017)

40 Deb, Donahue, and Glaisyer (2017)

reason but instead on emotion and bias.⁴¹ With such a pessimistic view of the public's judgment, it is little wonder Progressives would seek to place democracy on what they perceive to be a surer foundation. Progressives' embrace of deferential democracy is also consistent with their *institutionalized* view of reliable knowledge acquisition. Like Lippmann and many of his contemporaries,⁴² Progressives today view knowledge as best acquired by *experts* working within *institutions* that have developed *reliable methods* for discovering and disseminating knowledge. The most notable similarity that ties together the “institutions of education, science, and media” mentioned by Schiffman as having “traditionally served to keep 'false facts' and demagoguery at bay” and favored by many other Progressives is that each of these *is* an institution, one dedicated to the production and dissemination of knowledge and whose members are trained in a certain canon of *established methods*. While, as I have already emphasized, Progressives do not usually feel the need to explain *why* they favor the institutions they favor, those few who do often cite the methods employed by these institutions, such as scientific peer-review and journalists' methods for choosing and naming sources, as a primary reason.⁴³ Since Progressives seem clearly to view the knowledge acquired and disseminated by these methods to be highly reliable whereas the judgments of ordinary people consulting their own experience are much less so, it makes sense that they would see routinely deferring to the claims of experts trained in such methods as a better foundation for civic epistemology than (what they take to be) citizens' highly unreliable judgments.

Points of Agreement with Progressives

The reader will by now have probably guessed that I do not agree with either the Progressive point of view toward post-truth or with deferential democracy as a good model of civic epistemology. That is true, and much of the rest of both this introduction and the dissertation as a whole will be dedicated to explaining why. However, before I launch into my disagreements with Progressives and the deferential democracy they endorse, I want to stress the

41 Dahlgren (2018), Hannan (2017)

42 The intelligentsia of Lippmann's generation of the 30s and 40s, writes Schudson, was rife with “distrust of the public and a doubt that representative institutions could ever act wisely,” prompting them to embrace professionalism whose members developed a “proprietary attitude toward 'reason' and a paternalistic attitude toward the public” (Schudson 1978, 127-129).

43 E.g. Goldman (2010)

points on which I agree with them. First, I agree with some of their pessimism about the judgment of the lay public. Members of the lay public are not experts in all or even most of the epistemic domains that bear implications for public policy. Because of this, in very many cases we should naturally expect their judgment to be much poorer than that of experts. Members of the lay public are also plagued by a multitude of cognitive biases that skew their judgments, and in many cases our current institutions of democratic decision-making provide few if any strong incentives to motivate them to combat those biases. In both these ways, I agree to a certain extent with Progressives who, as we shall see, heavily emphasize citizens' reliance on experts and their susceptibility to cognitive biases.

In addition to sharing some of their pessimism about lay citizens in the contemporary context, I agree in large part with Progressives about the importance of knowledge-producing institutions, like those of the “scientific community,” whose members are trained in established research methodologies and whose knowledge-production process adheres to them. We owe such institutions a tremendous debt of gratitude for their contribution to our current stock of knowledge and to many of the technologies that make our lives healthier, safer, and more enjoyable than ever before. Moreover, I tend to agree with Progressives' views of the *particular* institutions they trust and the *specific* beliefs they endorse, vis-a-vis those they identify with the post-truth movement. I believe, for instance, that the world is roughly spherical, that climate change is largely human-caused and likely to have detrimental climatological effects in the foreseeable future, and that vaccines are an extraordinarily beneficial and important part of public health. In these, as well as what I take to be the overwhelming majority of other cases, my beliefs coincide with Progressives, largely because I share their view that the experts within the relevant domains are both credible and trustworthy. Though I shall be spending a great deal of time in this dissertation disagreeing with and criticizing Progressives, I want to be clear that it is not their *specific conclusions* about *which* institutions and beliefs are worth embracing that is the primary focus of my criticisms. Rather, my primary objections in this dissertation are to Progressives' embrace of *deferential democracy* as a normatively desirable model of civic epistemology and the *deep pessimism about lay citizens' judgment* on which it is founded.

Problems with Progressivism

These points of agreement notwithstanding, in this dissertation, I will argue that there are two major problems with the Progressive view. First, it is *descriptively misguided*. By this I mean the Progressive view (a) is *excessively pessimistic* about the judgment of lay citizens; (b) is *excessively credulous toward* established epistemic authorities; and (c) implies a naive, or at least dangerously short-sighted, view of the *likely outcomes* of citizens adopting its preferred deferential attitude toward such authorities over time. Progressives' deep pessimism about lay citizens is based on an exaggerated view of the role and immutability of cognitive biases in human reasoning and a tendency to fixate on sensational but not necessarily representative examples of epistemic behavior online and to ignore more positive examples. It is also jarringly at odds with the performance of liberal democracy for the past 250 years, which, I will argue, is fully inexplicable if citizens' judgment is as hopelessly poor as Progressives often portray it to be. At the same time, Progressives' credulity toward established epistemic authorities is excessive. This is *not* because the epistemic authorities we inherit are *typically* wrong or untrustworthy. On the contrary, I believe that it is generally the case that the authorities we inherit *are* reliable and that this is very likely the case with most of Progressives' favored epistemic authorities. Nevertheless, Progressives' credulity toward established epistemic authorities is excessive because even if these authorities are *typically* trustworthy *right now*, adopting an attitude of *routine deference* to their claims actively discourages citizens from discovering their errors made in good faith, uncovering corruption, and innovating new and better ways of thinking than those currently practiced by the accepted epistemic authorities. In other words, Progressives' advocacy of routine deference to established authorities and their tendency to excoriate anyone who does not place their trust in those authorities favored by Progressives actively disincentivizes citizens from using the best tools available to us for *improving our collective knowledge* and *discovering and resisting illegitimate authority*. The likely outcome of the adoption of such an attitude toward established epistemic authorities in the long term is for the authorities themselves to become either *complacent* or *actively abusive* due to the easy and predictable quiescence of the public. Thus, even though I agree with Progressives about many of the specific beliefs and knowledge-acquiring institutions they embrace, I believe

that as a *descriptive* matter we have strong reasons to object to the deferential ideal of civic epistemology they endorse.

The second problem with the Progressive view, I will argue, is that it is *incompatible with the ideals of liberal democracy*, which is the form of democracy many Progressives see themselves as defending and the only form of democracy I view as *worth* defending. Liberal democracy's support for a combination of constitutionally protected freedoms of communication and universal, equal suffrage make no sense unless we think citizens ought to base their beliefs and decisions on their *own judgment*, but that is precisely what is *discouraged*, even *disparaged*, by Progressives. Progressives are also out of step with the ideals of liberal democracy many of them claim to embrace with regard to their views of the proper role and meaning of “freedom of the press,” of which many of them claim to be staunch advocates. As we will see, for Progressives, “the press” means *the legacy news media*, the purpose of “freedom of the press” is to enable the specialized group of professionally trained journalists employed by the legacy news media do their job as society's “watchdogs,” or “the fourth estate,” keeping the public abreast of important affairs and ferreting out corruption among political leaders. Likewise, “freedom of the press” is viewed by Progressives as being *threatened by* practically anything that undermines the legacy news media, including declining public trust, the words of political leaders, and the rise of digital-era competitors. In the concluding chapter of this dissertation, I will argue that these attitudes stray very far from the ideal of “freedom of the press” liberalism originally – and rightly – endorsed, both in terms of *meaning* and *function*, and that digital-era platforms like Reddit, Twitter, Youtube and Facebook more closely reflect that meaning and are in many – but not all – ways better equipped to perform that function than the legacy news media.

Like Progressives, I view liberal democracy as worth defending. I also share their opinion that the judgment of lay citizens is often inadequate and/or flawed, and that many of our established institutions for acquiring knowledge are largely reliable and that many of their claims consequently deserve to be trusted. However, I disagree with Progressives' pessimistic interpretation of the meaning and impact of 21st century communications technologies on democracy. I also disagree with their response to those technologies, both in terms of the public policy approach they advocate with regard to them and in terms of the attitude of suspicion and derision they show toward them. Progressives' response to our changing 21st century

communications landscape implies an attitude of deep pessimism about the judgment of the lay public and a deferential model of civic epistemology, both of which are descriptively misguided and out of line with liberal democratic ideals.

Doxastic Self-Rule as a Core Component of Civic Epistemology for Liberal Democracies

In the real world, all serious thought is comparative. By this I mean that, since there are limitations in our understanding of the world and our execution of plans, political or otherwise, merely identifying a series of flaws in a given way of thinking about or dealing with an issue, even serious ones, is not sufficient to demonstrate that we oughtn't think that way. Instead, one must identify an alternative way of thinking that is superior to the one whose flaws are identified.⁴⁴ Identifying the problems inherent in Progressive view is beneficial only insofar as we can conceive of another, better way of thinking that is not subject to those problems and does not, in solving or avoiding them, create worse ones.

I believe that not only is there a better model for conceptualizing the relationship between the judgment of lay citizens, deference to authority, and the well-being of democracy than the one subscribed to by Progressives, but that the very post-truth movement condemned so vociferously by the Progressives can serve as a basis for such a model. I believe Progressives do a deep disservice to their democratic co-citizens by characterizing the post-truth movement as a descent into some permutation of tribalism, nihilism, solipsism, irrationality and/or other epistemic vice. A better interpretation is that of sociologist Steve Fuller, who describes the post-truth movement as an attempt by such citizens to *take responsibility* for the decisions they are required to make by virtue of being citizens in a democratic society.⁴⁵ Contrary to the caricatures painted by Progressives, most post-truthers do not reject the very *idea* of truth or “facts.” They do not even reject the idea of *authority* or its fellow traveler *expertise*; instead, it is the expectation of *routine* deference to *particular* established authorities on the basis of their status *as* established authorities that they reject. Rather than simply absorbing accepted views about who the “authorities” are and deferring to them as a matter of course, post-truthers insist upon

⁴⁴ This principle, commonly acknowledged by political thinkers, applies far beyond the gritty, realistic world of politics. As Imre Lakatos has argued, scientific knowledge is subject to the same limitations that curtail our search for perfect answers in the realm of politics and so must be measured against the same comparative standard (Lakatos 1978).

⁴⁵ Fuller (2018), 78

taking responsibility for such deference, and by extension the beliefs formed and decisions made on its basis, by requiring that it be justified by the lights of their *own* judgment.

The signature feature of post-truth is, I will argue, not the abdication but the *assumption* of the epistemic burdens that go hand-in-hand with liberal democratic citizenship. I call this assumption the exercise of *doxastic self-rule*. The essence of *doxastic self-rule* is the insistence that the beliefs one adopts ought to be *justified by the lights of one's own judgment*. The desire for citizens to exercise doxastic self-rule in this way is, I will argue, an essential component of any model of civic epistemology that wishes to be consistent with the ideals of liberal democracy. According to doxastic self-rule, citizens ought to *judge for themselves* which claims are true, which claimants ought to be trusted, and which public policies are justified. Their basis for adopting and/or communicating a given belief ought to be its persuasiveness as assessed by the lights of their own judgment, not because of reputation held by given claimant, or the possession of any other marker of social position or signifier of prestige attributable to that claimant. On this view, *no* deference to *any* epistemic authority is justified by any other means than a *judgment* on the part of the citizen that such an authority both (a) ought to count as a *legitimate* authority in that epistemic domain (aka an “expert”) and (b) is likely to be a *trustworthy source* of information given the situation at hand. Again, this line of thinking should be friendly to proponents of liberal democracy. A key point of emphasis for liberal political theorists in recent decades has been that *authority* needs to be *justified* in order for demands of *deference* to it to count as *legitimate*.⁴⁶ I believe a similar view is implied by most of our views of democracy, as embodied in slogans such as “no taxation without representation.”

The insistence of post-truthers that deference to authority can be justified *only* by an active exercise of judgment on the part of the citizen is, I believe, far more in line with the ideals of any democracy worth defending than the insistence upon rote deference to pre-defined authorities encouraged by Progressives. It is more desirable both in terms of its compatibility with the notions of self-rule endemic to any desirable version of liberal democracy and in terms of its facility in fostering good political outcomes, such as deterring complacency and abuse by established authorities and facilitating the growth of collective knowledge over time. For these

46 The demand that authority must be justified to those over which it is wielded has been emphasized by liberals in recent decades (e.g. Waldron 1987, Rawls 2005) to such a degree that, in the words of Simone Chambers, “justification stands at the center of contemporary liberal theory” (Chambers 2010, 893).

reasons, I will argue that a *post-truth citizenship* in which citizens insist on predicating deference to epistemic authorities on their *own judgment* is a better model of civic epistemology for any version of democracy worth defending than the model of *democratic deference* offered by Progressives. Despite my agreement with Progressives on a number of issues in the contemporary political landscape, then, my primary aims in this dissertation will be to argue against the model of civic epistemology they endorse and in favor of one closer to that embodied by the post-truthers they oppose.

Outline of the Dissertation

Each chapter of this dissertation will show evidence demonstrating the deep pessimism about lay citizens' judgment and the attitude of routine deference to established authorities encouraged by the Progressive response to the rapidly changing informational environment of the 21st century. In each chapter, I will point out why I think this pessimism is unfounded or excessive and why this attitude of routine deference to established authorities ought to be jettisoned in favor of encouraging citizens to assertively exercise doxastic self-rule in deciding whom to trust, which specific beliefs to endorse, and whether and how to pass on the beliefs they hold to their peers. The first chapter will focus on the concept of *liberal democracy*. In it, I will explain what I mean by the term “liberal democracy” and why I, like many Progressives themselves, think it is the sort of democracy most worth defending. While I share this belief with many Progressives, I will argue that the deferential democracy they endorse is at odds with the core ideals of the very liberal democracy they see themselves as defending. Specifically, it is at odds with liberal democracy's *strong preference for doxastic self-rule*. I will argue that liberalism's rights of communication and democracy's support for universal, equal opportunity of suffrage make no sense – indeed, seem actively destructive, or at least deeply reckless – unless we carry a *modest optimism* about the judgment of the lay public and see citizenship as properly practiced by citizens assertively *using* that judgment to decide for themselves which authorities to trust and which beliefs to adopt. I will also argue that both liberalism and democracy, rightly conceived, require *all authority to be justified by reason to those citizens over which it is to be exercised*, and insist, against certain theorists in the “public reason” tradition of liberalism, that the criterion that decides whether that authority *is* justified cannot be specified in advance by the

theorist but is instead each citizen's own *judgment*. I will argue that all these features – modest optimism about lay citizen's judgment, the encouragement of doxastic self-rule, and the insistence that authority may only be justified by an appeal to the judgment of lay citizens – are more characteristic of the post-truthers Progressives deride than of Progressives themselves, and so more in line with the ideals of liberal democracy.

In Chapter Two, I will zoom out and explore an unanswered question I take to be implicit in the myriad discussions about our evolving communications landscape, and its effects on democracy, that preoccupy so many concerned commentators at present. This is the question, “What does it *mean* for democracy to be doing well epistemically?” I will argue that while Progressives commonly make assertions about the negative – indeed, democracy-threatening – epistemic impact of 21st century communications technologies, they do so without first clearly specifying what democracy would look like if it was *doing well* epistemically. I will do that work for them, deriving a model of democratic epistemic well-being from what I take to be a charitable interpretation of Progressives' own work. I will then derive clear criteria by which the state of contemporary democracy can be assessed according to this model. I will argue that Progressives' assertions about the epistemic state of democracy never meet all of these criteria and that, worse, in many cases close attention to the criteria suggests that the *same evidence* cited by Progressives as evidence of democracy's epistemic peril might just as well be interpreted in the *exact opposite* way, that is, as a sign of democracy's sound and/or improving epistemic well-being. I will show that Progressives' tendency to draw pessimistic rather than optimistic *conclusions* about the epistemic impact of 21st century communications technologies on democracy often stems from their having *presupposed* a deeply pessimistic view of the judgment of lay citizens, and one that is not clearly justified. I will conclude the chapter by connecting Progressives' encouragement of a civic epistemology of routine deference to established epistemic authorities to their unjustified pessimism about lay citizens and the unjustifiedly pessimistic view of the epistemic impact of 21st century communications technologies built upon it. Since Progressives' deep pessimism about lay citizens is unjustified, I will argue, so too is their emphasis on the democratic importance of citizens *routinely* deferring to *established* epistemic authorities rather than basing their decisions of which authorities to defer to on an exercise of their *own judgment*. This is a good thing for liberal democracy since, as argued in

Chapter One, liberal democracy makes sense only if one carries a modestly optimistic view of lay citizens' judgment and encourages those citizens accordingly to exercise doxastic self-rule. However, it *does* begin to call into question the extent to which Progressives can sensibly and coherently claim themselves to be proponents of liberal democracy. Throughout the rest of the dissertation, I will build the case that they cannot, and that Progressives have generally responded to this dilemma by advocating approaches to contemporary political issues and views of citizenship that are out of line with the liberal component of liberal democracy. I will argue that we should not follow them in doing so, partly for the reasons spelled out in Chapter Two.

In Chapter Three, I will continue the discussion about what it means for democracy to be doing well epistemically in another way. This time, I will do so by examining the relationship between the behaviors of *individual* citizens and the *collective* outcomes likely to result from those behaviors. I will identify three contemporary lines of conversation in which Progressives commonly cite common sub-optimal epistemic behaviors of citizens, or what I will call *epistemic vices*, as evidence of the poor or declining epistemic well-being of democracy as a whole. I will argue that this approach is mistaken. Using a variety of historical and contemporary examples and theoretical arguments, I will argue that we cannot straightforwardly extrapolate the epistemic well-being of democracy *as a whole* from the individual epistemic habits and behaviors we observe among lay citizens at a given moment; on the contrary, in many cases, epistemic behaviors that lead individual citizens or even large groups of them to incorrect conclusions in a given case are actually signs that the conditions most likely to foster the overall epistemic well-being of democracy. I will argue that this is precisely the case for many of the specific epistemic vices fixated on and treated as hallmarks of the post-truth “malaise” by Progressives. A common thread in all these cases will be that the epistemic behaviors that are rightly dubbed vices at the level of individual knowledge acquisition yet turn into virtues at the level of the collective are all the result of citizens assertively exercising doxastic self-rule. Were the individuals in these examples to be more timid and deferential in the way encouraged by Progressives, they would *individually* have increased their odds of adopting a higher percentage of true beliefs but this would have come at the *collective* cost of stifling the growth of knowledge in the long run. I will argue that this suggests we ought to encourage citizens to err on the side of assertiveness rather than deference when it comes to exercising doxastic self-rule in their

decisions about what to believe and whether and how to pass on those beliefs, since excessive assertiveness often leads to beneficial collective epistemic outcomes even if leads to detrimental outcomes at the level of the individual.

Whereas Chapters Two and Three took a macro-level view of the epistemic well-being of democracy as their starting point, the next two chapters will start at the level of the individual. In Chapter Four, I will take an in-depth look at the epistemic plight faced by individual citizens trying to decide what they should believe in today's enormous, complex and interconnected world and develop some general normative principles for belief formation under those conditions. I will argue that the case for adopting an attitude of assertive *doxastic self-rule*, i.e. insisting that claims pass a test of scrutiny in the light of one's own judgment if they are to be accepted as beliefs, is, if anything, even *more* compelling from the perspective of the individual than it is from the standpoint of the collective. Indeed, I will argue that doxastic self-rule must be the starting point of *any* model of normative epistemology if that model is to be persuasive because doxastic self-rule is the only means by which *both* individuals may consistently and intentionally *improve their beliefs* and *discover and resist illegitimate authority*. Likewise, only societies in which *individuals* practice doxastic self-rule have any hope of improving their beliefs and discovering and resisting illegitimate authority at the level of the *collective*. I will then apply this standard to the specific topic of *trust in testimony*, which is central to many discussions about citizenship in 21st century democracies. I will argue that the desirability for individuals to exercise doxastic self-rule applies to decisions about which testifiers and testimonies to trust just as much as it does to all decisions about what to believe, and for the same reasons. Thus, any convincing model of *civic epistemology* ought to encourage *epistemic vigilance*, or the regular, active exercise of doxastic self-rule by lay citizens toward the testimonial claims they encounter, *regardless* of whether or not those claims disseminate from a source that enjoys a reputation of authority or expertise. Progressives' model of civic epistemology does not meet this standard. Instead, it advocates an attitude of *selective skepticism*, encouraging citizens to be excessively skeptical of certain pre-defined types and sources of testimony, such as rumors and claims shared on social media, while simultaneously encouraging them to be routinely deferential toward Progressives' favored sources, such as the legacy news media. Worse, many Progressives support censorship of information online by some set of authorities, be it private or public. Because both

this disposition of routine deference to established authorities and censorship by their very nature actively inhibit the effective exercise of doxastic self-rule by citizens, I will argue that the Progressive approach undermines the sort of oversight that is central to the proper functioning of both liberalism and democracy, and as such are bound to eventually lead to complacency and/or abuse by those authorities. Rather than disparaging lay citizens' judgment in general, encouraging them to exercise routine deference to established authorities, and pushing for authorities to exercise more control over the informational environment, we ought to be *encouraging* citizens to exercise epistemic vigilance, *assertively* using their own judgment to assess the claims and claimants they encounter online, and insisting on being allowed to access and pass on information freely so as to be able to judge its merit for themselves.

In Chapter Five, I will address the issue of *trust in testimony*, whose role in both normative epistemology and in liberal democratic society have been the subject of heated debate in recent years. I will start by examining the epistemic side of the issue. After establishing that trust in testimony is a critical element in the adoption of the overwhelming majority of our beliefs and introducing some core concepts by discussing the debate between reductionists and non-reductionists in the epistemology of testimony, I will develop a normative approach to the question of whether and how far to believe testimony based on the *degree of consilience*, or convergence on a given hypothesis from multiple independent grounds, available to the would-be knower. *Ceteris paribus*, I will argue, the greater the degree of consilience, the more confidence one is justified in attaching to a given testimony or testifier. Given the enormous size and complexity of the contemporary political landscape, however, I will argue that the degree of consilience available to *anybody* trying to decide which testifiers and testimonies to trust under contemporary conditions is *limited* to such a degree that we ought not place too profound a degree of confidence in *any* single would-be authority's opinion about which testifiers and testimonies ought to be trusted. The idea that any authority exists who is bound to be better at deciding which news is reliable and/or which news sources ought to be trusted assumes a sort of expertise – expertise about the news – that I argue we have no good reason to suppose *anyone* to possess, and even if some entity *did* possess such expertise, it could not be demonstrated to citizens in a way that ought to be persuasive to most citizens. Moreover, even if all these things were not true, citizens would still have strong common-sense reasons to oppose the civic

epistemology of routine deference to established authorities and/or the censorship of news on the basis of fact, both of which positions, as we have seen, are increasingly advocated by Progressives. For the overwhelming majority of the news, lay citizens have no good reason to suppose that Progressives, the authorities they favor, or any single identifiable group is consistently better positioned to who tell *which* news is reliable or which news sources ought to be trusted than lay citizens themselves are. I will then explore a few of the most useful epistemic tools available to *anyone* attempting to decide which testifiers and testimonies to trust in contemporary conditions. I will find that those tools are *useful* in the sense that they can potentially be put to quite good effect in helping knowledge-seekers adopt reliable beliefs and avoid unreliable ones, but that because most of the phenomena reported by those testimonies are distant and unfamiliar, there is an essential *precariousness* to the beliefs we adopt solely on their basis. While this might seem to justify Progressives' deep pessimism about lay citizens' judgment, I will argue that it much more profoundly undermines Progressives' attitude of credulity toward “fact-checkers” and their easy disparagement of fellow citizens' conclusions about which testifiers and testimonies to trust. The most important conclusions from these considerations, I will argue, are that (a) there are no identifiable “experts” – that is, no one whose judgment is regularly and demonstrably better than that of lay citizens – when it comes to deciding which testimonies and testifiers to trust, and (b) while limited, the tools available to lay citizens for deciding which testimonies and testifiers to trust are still *useful*, especially when used with good *judgment*. This gives us another set of reasons to oppose the civic epistemology of routine deference to epistemic “authorities” embraced by Progressives and to opt for one that does not so easily fall prey to false notions of “expertise” and encourages citizens to more assertively exercise doxastic self-rule.

In Chapter Six, I will argue that the problems I have identified with the Progressive View stem largely from their embrace of an epistemic outlook I call *institutionalism*, which views objective knowledge as discovered by institutions while viewing the judgment of ordinary citizens as largely suspect. I will argue that this outlook overrates the reliability of the knowledge of Progressives' favored institutions, underrates the judgmental capacities of lay citizens, and encourages complacency and corruption in the institutions Progressives defend by impeding public oversight of those institutions. Rather than embracing this sort of institutionalism, I will

argue in favor of an epistemic outlook I call *epistemic pluralism*, which I argue informed the thinking of American founding fathers James Madison and Thomas Jefferson. Adopting this outlook can help us avoid the dangers of Progressives' institutionalism and recognize the considerable value to liberal democracy of the very 21st century technologies Progressives portray as threats to it. I will conclude by arguing against Progressives' growing embrace of censorship online, as well as the deferential model of civic epistemology they endorse.

Post-Truth Citizenship as a Model of Citizenship Suited to the 21st Century

My hope in this dissertation is to try to articulate and defend a nascent movement that is embodied in many scattered places throughout democratic societies today but which has yet to be expressed in a clear and coherent fashion. This movement offers those of us who live in the early 21st century a chance to respond to the difficult and at times bewildering problems of democratic citizenship in large, complex societies in a better way than those before us responded to the disillusionment brought on by the failures of Progressivism, and to make better use of the 21st century's unprecedentedly powerful tools than the reflexive opposition to and demonization of those tools, which is the norm of Progressives and, unfortunately, seems to be the modal approach in our era. For we live in our own era of disillusionment twice over, an odd sort of double-disillusionment made possible by the murkiness of our own sense of history. Somehow, we have managed to inherit *both* the optimism of Progressivism and its rosy-eyed view of the “omnicompetent citizen” *and* the pessimism of post-Progressives, and the deferential model of citizenship they came to embrace, all at once. Because that original Progressive view was unrealistic, we cannot help but be disappointed when we observe the actual behavior and decision-making of democratic citizens in the present day. And because of our inheritance of post-Progressivism's deferential model of citizenship, many of our first instinct upon being disillusioned with the public is to rush into the arms of the authorities that emerged in the post-Progressive world, authorities endorsed by post-Progressives in an attempt to save democracy from itself. Yet, that very rush is the source of a second disillusionment – a disillusionment with these authorities themselves. It is as obvious as any fact of political life can be – in fact, it is the dominant fact *of* contemporary political life – that we live in an age characterized by widespread disillusionment with established authority. Progressives tend to either write off this

disillusionment as irrational or to seek to design institutions less susceptible to the forms of degradation, decay and corruption over time than those designed by their forebears. Either way, their conception of democracy and their approach to democratic citizenship is essentially the same: Democracy's well-being depends critically on the willingness of the public to readily defer to some pre-defined set of “the right” authorities. Good citizenship is deferential citizenship. A healthy democracy is a deferential democracy.

I hold a different view. I believe contemporary citizens' disillusionment with the authorities we have received as our civic inheritance is not just *not* irrational but is in fact largely justified. It is justified *not* because the institutions post-truthers view with suspicion – entities and institutions of media, government, and “expertise” broadly defined – are populated by especially evil individuals or are especially poorly designed, but because *all* human-created institutions are subject to inevitable decay over time. This decay is the natural result of the conflicts of interest between the public and those who hold positions of authority, the latter of which always has, in the words of political theorist Evan Turner, “an interest in preserving the power that goes along with their status” that prevents them from ever quite being “positionally neutral.”⁴⁷ The single best deterrent we know of for combating such decay to whatever degree we can, one whose importance has been historically emphasized by liberal and democratic theorists alike, is by subjecting such authorities to *external oversight*. But this is precisely what is *discouraged* by the deferential attitude advocated by Progressives and the managerial approach to the 21st century informational environment they have increasingly come to embrace. Indeed, contemporary efforts by the public to hold would-be authorities to such a standard of accountability – accountability *to the public* – are routinely condemned as *threats to democracy*. Since I believe much of the current disillusionment with established authorities stems from the public's observation of those authorities' poor and declining performance over time, I believe the deferential citizenship recommended by Progressives to be not just misguided but actively destructive, as it seeks to foster the very conditions most likely to see that decay accelerate.

I believe there is a better way, one that does not fall prey to the self-deluding optimism of the original Progressives' ideal of “omnicompetent citizenship”⁴⁸ but which also declines to

47 Turner (2003), 55

48 Lippmann (1925), 21

swing too far in the opposite direction, advocating a disposition of docile quiescence in the citizenry that invites corruption and abuse by the authorities to whom they are asked to defer. I hope by the end of this dissertation to have made that position, already nascent in some of the words and deeds of the post-truth movement, clear. We have a chance to bequeath to our descendants an intellectual inheritance which, while as undoubtedly imperfect as any human ideal may be, is nevertheless better than the one we inherited ourselves. It is my hope that these words may contribute in some small way to that quintessentially human endeavor.

Chapter 1: A Democracy Worth Defending

In the introduction I said that many of those I have chosen to call “Progressives” avow support for what I call “liberal democracy,” but that this supposed support was strongly at odds with the public policy positions and notions of civic epistemology endorsed by these thinkers. The aim of this dissertation is to argue against these Progressive positions and in favor of public policy positions and notions of civic epistemology that are a better match for liberal democracy. One might, however, wonder whether it's really so important to cling to this thing called “liberal democracy,” rather than simply embrace the public policy positions and notions of civic epistemology favored by Progressives even if it *does* mean abandoning liberal democracy. To answer that question, I'll need to explain what liberal democracy is and why I think it is of such great importance. That is the task of this chapter, which will be written in two sections. In Section One I'll define what I mean by “liberal democracy” and offer a defense of what I call “minimally consequentialist liberal democracy,” which I view as the only sort of liberal democracy worth defending. In Section Two I'll argue that it makes no sense to defend this sort of liberal democracy unless one holds a *modestly optimistic* view of the judgmental capacities of lay citizens and encourages those citizens to engage in *doxastic self-rule*, both of which stand in stark contrast to the approach favored by Progressives. This argument will set the foundation for the arguments of later chapters, in which I will advocate holding true to the ideals of liberal democracy and resisting the approach favored by Progressives.

Section One: Liberal Democracy

Definition of Liberal Democracy

Both “liberalism” and “democracy” are essentially contested terms whose proper meaning and correct use have been responsible for the spilling of much academic ink. The situation regarding the term “liberal democracy” is not much better. While I believe this ink not

to have been spilled in vain (well, not *all* of it), it is not necessary for me to get caught up in such arguments to make the points I want to make in this dissertation. Rather than arguing about what “liberalism,” “democracy” and “liberal democracy” *really* mean, I will simply supply the reader with up-front notice about what *I* will mean by them.

When I use the term “liberal democracy,” I mean a *political system*, and the *philosophical defense* of that system, which, whatever else it entails, includes the two following components. First, it enshrines *constitutional protection* of rights intended to protect what I call *freedoms of communication*. The most important of these rights for my purposes are *freedom of speech* and *freedom of the press*, but freedom of religion and the freedom to assemble are also related to this canon of rights. This is the liberal component of “liberal democracy.”⁴⁹ The second core component of what I call “liberal democracy” – the democratic one – is *universal, equal opportunity of suffrage*. This means that *all* able-minded citizens who have reached their majority have an opportunity to take part in decisions about who is to hold political office and certain other constitutional or legislative decisions, such as amendments and referenda. I phrase this as equal *opportunity* of suffrage mainly to accommodate for various forms of democracy besides the form of popular voting we are all familiar with, for example by use of sampling or lotteries, which have experienced a resurgence of popularity among some democratic theorists in recent years⁵⁰ and to which I hold an open ear. When I say “liberal democracy,” then, I mean a *system of government that enshrines in its constitution (a) protections for core freedoms of communication, especially freedoms of speech and the press, and (b) universal, equal opportunity of suffrage among able-minded adult citizens.*

49 Note that this definition of “liberal” entails no specific claims about economic policy or private property. This is both because I see the most important legacy of liberalism for the purposes of this dissertation to stem from its defense of the rights mentioned above and because there is an enormous range of views among self-identified “liberals” who defend such rights with regard to economics and the proper definition and delineation of private property rights. My own position is that the idea of private property in the sense of a sphere of action that is beyond the jurisdiction of the public is indispensable both to the good life and to the other aims of liberalism. However, this still leaves much room for disagreement. One thing I want to be clear my endorsement of my embrace of this feature of liberalism does *not* necessarily imply – indeed I would argue it *ought not* imply – is that the activities of large corporations ought to enjoy the same immunity against public accountability as individual citizens. Regardless of this particular position, though, I just want to be clear that my definition of liberalism intentionally revolves around the sorts of civil rights alluded to in the text, and is amenable to a wide range of views about economic policy and private property.

50 E.g. Landemore (2012), Fishkin (2009)

My defense of liberal democracy in this section will take into account both ideal and practical considerations. Ideally, I will argue that the core freedoms of communication cherished by liberals and the equal opportunity of suffrage offered by democracy are partly constitutive of our views of “the good life,” and are therefore desirable in themselves. Practically, I will argue that liberalism's freedoms of communication are of extraordinary importance in pursuing two vital societal aims, the aim of *improving our collective knowledge* and the aim of *discovering and resisting illegitimate authority*, while democracy's equal opportunity of suffrage offers a clear and comparatively easy to implement means for citizens to take advantage of these benefits of liberalism politically. I will concede that the justification of liberal democracy rests crucially on a certain kind of optimism about lay citizens, particularly their capacities of judgment on matters of political importance, that might seem to pose a challenge to liberal democracy on the grounds of realism. However, I will argue that the historical performance of contemporary liberal democracy gives us no reason to suppose that citizens cannot be relied upon to meet these demands. I will conclude that, given its desirability in both ideal and practical terms and the absence of any compelling reason to suppose it is impracticable, any normatively compelling defense of democracy must be a liberal democracy.

Minimal Consequentialism

I mentioned above that political theorists have defended many different versions of liberalism, democracy and liberal democracy. Likewise, the arguments supplied in favor of each of these systems have also varied considerably. Take democracy. I have defined democracy as requiring that “all able-minded citizens who have reached their majority have an opportunity to take part in decisions about who is to hold political office and certain fundamental matters of constitutional law.” In contemporary times, the primary means through which this opportunity has been afforded has been the vote. Even if we politely agree not to quibble over this as an acceptable definition of “democracy,” we must still wrestle with the fact that political theorists have spent a lot of time arguing about just why (and to a lesser extent whether) this sort of rule by the people is a good way of selecting political leaders and/or public policies. The number of different views on the matter is so extensive as to threaten to throw the average democratic citizen, born and raised in an era where democracy's superiority is accepted with scarcely

questioned universality, into an existential panic. One major axis of disagreement between democratic theorists centers around whether democracy is ultimately justified by its *fidelity to procedure* or whether it is justified at least in part by the *quality of decision-making* that results from democratic procedures. Political philosopher Jamie Terrence Kelly does a nice job of parsing democratic theories along this axis. At one extreme are what Kelly calls “purely procedural” justifications of democracy. As Kelly points out, these theories of democracy “do not rely on any substantive claims about the correctness of political judgment: as long as these procedures are in place, democracy's value is assured, regardless of how individuals judge.”⁵¹ In contrast, all other democratic theories rely on at least a minimal claim about the tendency for democracies to produce good (enough) public policy decisions to justify preferring them over other available political systems. The threshold for what counts as “good (enough)” varies enormously between theorists, with optimists like Rousseau seeming to believe democracy properly practiced can do no wrong⁵² and less romantic theorists insisting that democracy is doing its job just by ensuring peaceful transitions of power along (perhaps) with moderate protections against tyranny.⁵³ There are, of course, also many other defenses of democracy that lie between these two extremes.

What matters for my purposes is not which specific theory of democracy one embraces but rather whether that theory is *purely procedural* or relies on at least *some* claim about the quality of decision-making that tends to be accomplished by democratic institutions and procedures. I will call this requirement that democracy (or any political system) produce good public policy outcomes in order to be justified *minimal consequentialism*. I am a *minimal consequentialist* in this way. I do not agree with the notion that any normative defense of any kind, be it a political system, moral framework, or who to play on this week's fantasy football team, could ever be properly justified without any reference whatsoever to its consequences. Accordingly, I reject pure procedural justifications.⁵⁴ Whatever the justification of democracy or

51 Kelly (2012), 46

52 This is but one of many possible interpretations of Rousseau's rather confusing discussion of “the general will” in *The Social Contract* (1952 [1762]).

53 This is my reading of Walter Lippmann's view in *The Phantom Public* (Lippmann 1925, ch. 4) and Adam Przeworski's “minimalist defense of democracy” (Przeworski 1999).

54 Some political theorists may wish for a more elaborate defense of my position here, but I don't think the choice between pure proceduralism and theories that include some consideration of consequences is the sort of thing about which decisive arguments can be made. Like my preference, all things being equal, for pleasure over pain,

any other political system, in order to be persuasive it must be defended at least partly because it tends to produce good public policy outcomes. This is exactly the case for my support of liberal democracy. I support liberal democracy because, in addition to being in harmony with certain notions of fairness and the good life I hold dear (to be elaborated further down), I believe it tends to produce good public policy outcomes, over the long run, to justify choosing it over the other options that seem available to humankind at present.⁵⁵ Thus, any version of liberal democracy I support must meet the standard of minimal consequentialism, which I define as follows:

Standard of Minimal Consequentialism: Any political system must be justified at least in part its tendency to produce good public policy outcomes, compared to other viable alternatives.

Primary Goods and Primary Bads as Benchmarks for Minimal Consequentialism

What, specifically, does minimal consequentialism require? I cannot articulate a comprehensive account here, but I believe useful benchmarks to be *primary goods* and *primary bads*. Minimal consequentialism as I envision it entails both *identifiable (if gradual) progress over time in attaining primary goods* and *far better than random tendency to avoid primary bads*.⁵⁶ Among the most important primary goods are the improvement of our collective knowledge, peace, stability of a kind that allows citizens to make plans in reasonable confidence that they can be carried through to completion, bodily health, material security such as ample food and shelter, and at least some appreciable freedom to pursue one's vision of "the good life." Among the most important primary bads are corrupt authorities, war, instability of a sort that undermines citizens' confidence that any plans they form have hope of being carried through to completion, hunger and poverty, widespread lack of access to health care, and oppression defined

it is a matter of what kind of world I would *prefer*, based on my own tastes and sensibilities, to live in. Since it is a "mere" matter of preference in this way, simply stating my preference for one thing over another is, I think, the appropriate way to justify my position.

55 Overall, the argument I'm defending here is similar in nature with that provided by David Estlund, who argues that democracy's legitimacy as a political system stems in part from its ability to produce good decisions over time and in part from its fidelity to other values we cherish (Estlund 2008, 167).

56 The latter of these two desiderata was inspired by Estlund (2008, 162), but the avoidance of such primary bads is the only requirement Estlund believes to be necessary in order to fulfill the consequentialist component of justified "democratic authority." I think Estlund is close to the mark, but that we should also insist on at least a modest tendency to *secure* primary goods, not just avoid the bads.

as the exercise of power in a way that denies to citizens their own vision of the good life along important axes such as choice of religion, occupation, location, etc.

I believe liberal democracy to be a more stable and reliable means of meeting these standards than other available alternatives. In part, this is because I believe conditions of liberal democracy to be *partially constitutive of* some of the desiderata listed above. For instance, I share the view of Walter Lippmann and Adam Przeworski that one of the chief benefits of popular elections is that they constitute “conflict without killing,”⁵⁷ a peaceful vehicle for exerting “in civil society...the force which resides in the weight of numbers.”⁵⁸ Insofar as citizens agree to decide who gets to hold power next by election rather than an “appeal to heaven,”⁵⁹ democracy actively encourages the maintenance of peace and discourages the onset of civil war. Beyond this, many citizens enjoy the very *fact of* having a say in fundamental political decisions that affect them and those they love. In these ways, living in a democracy is partially constitutive of the “good life.” Similarly, the suite of rights prioritized by liberalism enshrines many of the freedoms whose exercise is both itself partly constitutive of the “good life,” as well as being necessary in order for each citizen to pursue other aspects of whatever the “good life” means to them. For example, citizens exercise their freedom *in order to* practice their chosen religion, but additionally the *very experience of* exercising that freedom is part of many people's (and my) view of “the good life.” Insofar as liberalism and democracy are *partially constitutive of* certain primary goods, and the avoidance of certain primary bads, liberal democracy inherently meets some demands of minimal consequentialists that cannot be met by political systems that entail neither liberalism nor democracy.⁶⁰

57 Przeworski (1999, 49)

58 Lippmann (1925, 58)

59 This is of course the poetic term John Locke used for civil war in his famous *Second Treatise of Government* (Locke 1980 [1690], §168).

60 Of course, liberalism doesn't please everyone. It is especially upsetting to those who detest the points of view and/or lifestyles of individuals and cultures that differ significantly from their own yet are allowed to persist in a liberal society, as well as to people who wish for strong, state-backed assurance that their own lifestyle and traditions will persist in perpetuity. There are many who so strongly value these things that they would rather live in an illiberal society than give them up. All I can say in response is that I do not hold conservative values like this in high enough regard to view them as trumping the desirability of liberalism. Liberalism offers a wide range of lifestyles the opportunity to perpetuate themselves through persuasion, enculturation, childhood upbringing, and many other means other than force of law. It also offers considerable sway for people to associate with whoever they please and refuse to associate with those they do not. In so doing, liberalism shows a certain degree of respect even for those with a conservative or illiberal temperament. However, I believe liberalism necessarily entails a *positive endorsement* of the idea that, in general, facilitating diversity and tolerating difference are normatively preferable to furthering what Rawls called particular “comprehensive doctrines”

However, I believe the minimal consequentialist's demands that a political system make identifiable (if gradual) progress over time in attaining primary goods and exhibit a far greater than random tendency to avoid primary bads are fulfilled by liberal democracy even beyond situations in which liberal democracy is partially constitutive of those outcomes. Whether this is the case is, of course, at root an empirical question. Unfortunately, it is an issue upon which I believe we are very far from having rigorous empirical proof or even knowing how to go about getting such proof. Not only are “liberalism” and “democracy” difficult concepts to operationalize and measure in ways that seem both theoretically sound and relatively straightforward to observe, but they co-vary with too many other phenomena – military power, population, land area, the legacy of colonialism, natural resources, economics, etc. – to demonstrate on empirical grounds what contribution the “liberal” or “democratic” part of the equation is making to the outcomes one observes. In the absence of any hope (at least at present) of rigorous empirical proof of their connection to good outcomes, I will have to rest my case for democracy's and liberalism's contribution to identifiable (if gradual) progress over time in attaining primary goods and far better than random avoidance of primary bads on a combination of theory and what I take to be historical evidence in support of that theory.

Let us begin with liberalism. I believe liberalism's suite of rights contributes to the realization of primary goods by encouraging experimentation along a great many axes – artistic, scientific, commercial, etc. – in a way that spurs both the *growth of our collective knowledge* and innovation on all these fronts. This knowledge enables us to act more effectively in the world by more reliably connecting means to ends, while these innovations produce a regular proliferation of benefits to the public in the form of aesthetics, cheaper and more useful products, advances in medical technology, and so forth, all of which have positive implications for liberal societies' ability to secure primary goods.⁶¹ As for the growth of our collective knowledge, the theoretical case connecting liberalism to collective epistemic improvement has been forwarded by countless

(Rawls 2005).

61 It may be worth reiterating my intention to refrain from explicit commitment to any particular views about the proper definition and scope of private property beyond the general desirability of having some appreciable space carved out in that society in which citizens retain qualified immunity from public censure, which I believe to be the essence of the term “private sphere.” I believe the commercial benefits mentioned here are likely to arise in any society that incorporates some system of freedom of exchange in which that kind of qualified immunity is enjoyed to some degree by market actors, and so my claims do not depend on any particularly strong endorsement of *laissez-faire* economics.

theorists and is, I believe, very strong.⁶² It is both rooted in a sound grasp of human epistemology and consistent with the intellectual and technological progress that has erupted in the West in the past half-century, whose onset coincides with the Protestant Reformation, the invention of the printing press, and the development of science, all of which represent the implementation of key aspects of liberalism in Western societies throughout that period. Likewise, liberalism's enshrinement of freedoms of speech and press helps to avoid primary bads. One of the main ways it does so is by making it easier for citizens to *discover and resist illegitimate authority*. It helps maximize society's ability to uncover conspiracies, malfeasance, deception, and other abuses of power, both by those in political office and by those in private stations that nevertheless heavily impact the public. Again, our reasons for supposing liberalism to be a good means of helping citizens discover and resist illegitimate authority stems both from a sound grasp of certain fundamentals of human psychology and sociology, as well as observed historical trends. A sound grasp of psychology and sociology tell us we should always expect there to be plenty of people out there seeking to gain power in order to advance their own interests at the expense of the public, while a sound grasp of sociology means that we should always anticipate pressure, both in terms of punishment from power holders and demands to conform to the existing order, to play a significant role in deterring those who might seek to question or contradict those who currently hold positions of authority from doing so. Given the public's interest in not being subjected to abuse, society needs to build in strong measures to counter these natural pressures. I believe the best way of doing so is by committing to protect the freedoms of communication traditionally defended by liberals. The reason *freedom* of communication is the best policy for facilitating these aims, rather than appointing any authority to maximize the *quality* of the information made available to citizens on the informational environment, is that *all* authorities have a conflict of interest when it comes to information that might facilitate the discovery of corruption by them and those they may be in league with. For this reason, the most reliable means available to the public for consistently facilitating the discovery of and resistance to illegitimate authority is to maintain communicative *freedom* rather than appoint some authority to try to maximize its *quality*. Again, the historical record supports these suppositions. The regimes around the world, past and present, that have the *least* such

62 E.g. Smith (1937 [1776]), Mill (2007 [1859]), Hayek (2011 [1960])

freedoms – prominently in our time, China, North Korea, and Russia, but by no means limited too these examples – usually exhibit an overwhelming tendency to manifest high levels of corruption and for those in power to abuse it with impunity. While the regimes that do the most to protect such freedoms of communication, like the Netherlands and the United States,⁶³ are by no means free from *all* such corruption, no serious thinker can suppose that power holders in these countries can abuse their office with anywhere near the degree of impunity enjoyed by those in the previous three countries mentioned or those others which are, like them, largely absent of liberalism's freedoms of communication. I believe a main reason for this is the strong liberal tradition of freedom of communication in the United States, the Netherlands, and countries like them. For all these reasons, I believe that liberalism is a strong impediment to primary bads.

What about democracy? I personally suspect it, too, contributes to the noticeable (if gradual) advance overtime in the attainment of primary goods for the people who live under that system of government, but I also believe the case to be mustered in support of this view is more tenuous than the case for liberalism, both in terms of its grounding in human psychology and sociology and in terms of historical trends. While it seems intuitive to expect better public policies to be sanctioned by leaders answerable to the public than those not, that tendency depends critically on the public maintaining an awareness and understanding of public policies and the behaviors of policy makers that is at odds both with the empirical record⁶⁴ and, in my view, realistic expectations of citizens who live in a complex world where their vote is typically but one among millions.⁶⁵ Because of the difficulty of keeping up with the complexities of

63 I use these two examples because they are the countries I am most familiar with that have the strongest protections of free speech, specifically. Many other countries who claim to endorse “free speech,” such as Great Britain and France, actually have extensive restrictions of and punishments for what the authorities in those countries deem “hate speech” as well, unfortunately, as increasing restrictions on information deemed by those authorities to be both factually false and dangerous to the public. I take these countries' regulatory approach to communications – and, by extension, that typical of most countries around the world – to be deeply in tension with liberalism and certainly, whenever that approach includes authorizing authorities to demand prior restraint, to be definitively anti-liberal.

64 See, e.g., Delli Carpini and Keeter (1996) and Somin (2016).

65 This case has been made many times over the course of the last century. The power of this line of argument connecting the consistently poor empirical performance of citizens when it comes to following and understanding politics to the lack of incentives faced by voters in large, complex societies can be inferred from its support by a very wide range of theorists who differ from one another in virtually every other respect. James Fishkin cites it in support of his preference for “microcosmic” deliberative democracy (Fishkin 2009); Richard Posner cites it in support of the competitive form of democracy he takes to be the status quo (Posner 2003); Walter Lippmann uses it to argue that our support for democracy shouldn't be related to good public policy

contemporary politics and the infinitesimal weight each individual vote has on influencing public policy, it should not surprise us to find, in the words of Przeworski, that “choosing rulers by elections does not [demonstrably] assure either rationality, or representation, or equality,” nor better economic performance, nor perhaps even “implementation of any criteria of justice,” compared with alternative forms of government.⁶⁶ I will not make any strong claims, then, about the tendency of democracy to facilitate the realization of primary goods. I will simply say that democracy does not appear to be demonstrably *more opposed* to that realization compared to non-democracies, and that my own suspicion, for whatever it's worth, is that in the long run democracy is indeed positively related to that realization.

The case for democracy when it comes to *avoiding primary bads*, however, I believe to be much stronger. While citizens may not always be knowledgeable enough about the details of public policy and current events to consistently select excellent leaders and back effective public policies, they *are* attentive to their own lives and those of the people they are close to, and many maintain a general low-level awareness beyond that circle as well.⁶⁷ When these groups' lives are subjected to significant setbacks in terms of rights, material well-being or general quality of life, the public can be counted on to notice such drastic changes and to punish those they perceive to be in charge for poor performance or abuse of office. In many such cases, unfortunately, the same complexities cited above make it difficult for the public to accurately place blame and identify preferable alternatives. This is perhaps most consistently evidenced by the tendency of democratic polities with a presidential system to punish incumbents in that office during times of economic downturn,⁶⁸ even though the causal effect on economic performance attributable to the head of the national executive office seems in most cases to be both limited and difficult to specify. Yet, I believe democracy deters primary bads even despite these epistemic difficulties. One reason is that even when the causal chain between policy and outcome is unclear, though,

outcomes (Lippmann 1925); and Jason Brennan cites it as part of his reason for rejecting democracy in favor of “epistocracy” (Brennan 2016). These examples far from exhaust the list.

66 Przeworski (1999, 43-44)

67 For example, scholars studying “economic voting” usually find that when it comes to the question of whether people vote on the basis of their own “pocketbook” or their perception of the overall national economic situation, the latter wins out (Stegmaier et al 2016, 587).

68 Stegmaier and colleagues argue that the existing literature suggests voters may tend to reward executive heads for economic performance just as much as they punish them for poor performance (Ibid., 587). Neither situation seem to paint an especially rosy view of the connection between voting and sound public policy, however.

the knowledge *that* noticeable decreases in citizens' quality of life are likely to be interpreted as reflecting poorly on their performance should exert a general pressure on incumbent power holders to do what they can to avoid the worst sorts of outcomes. Another is that many primary bads are in fact quite readily and reasonably attributable to specific policies and policy decisions than the general state of the economy. One of the most important primary bads for which this is the category of *corruption among authorities*. Gross violations of cherished freedoms, the imprisonment of citizens or the seizing of their property without due process, the “disappearing” of those opposed to the reigning regime on trumped-up charges, the regular demand of bribes in exchange for political services, and many other abuses meted out on citizens by political regimes with depressing regularity in the history of human affairs are easy to notice, quick to alarm citizens, and often clearly attributable to the authorities responsible for implementing them. If such policies are *approved of* by the public, perhaps because of widespread racism or some other prejudice, then of course democracy is no protection. But in such tragic cases, what *is*? In the many cases where the moral depravity of the public is not so dire, however, answerability of public figures to citizens constitutes a standing barrier discouraging would-be power abusers from causing such primary bads. The upshot is that, as legal theorist Richard Posner puts it, democracy serves “to keep the representatives on a tether, though a long one.”⁶⁹ That tether may not be as tight as we wish it could be, but it does much to protect the public from many of the primary bads most worth avoiding. In the language I will most frequently use in this dissertation, it greatly contributes to making sure the sort of authority employed in a democratic society is *legitimate*, that is, in keeping with appropriate normative standards. I believe democracy's facility in avoiding the proliferation of illegitimate authority and other primary bads makes it well worth appreciating.

Thus, my support for both liberalism and democracy stems at least partially from my belief that they facilitate good political outcomes compared to the alternatives at hand. Liberalism actively contributes to the realization of primary goods and the avoidance of primary bads, while democracy's most demonstrable value comes from its tendency to help us avoid the worst of the primary bads. But I believe the case for *liberal democracy* to be stronger than the

69 Posner (2003, 157)

case for *just* liberalism or *just* democracy.⁷⁰ Liberalism *facilitates* democracy's strengths and *mitigates* its vulnerabilities. It facilitates democracy's ability to fight illegitimate authority by creating the best-possible conditions for transparency about the mechanisms of those in power. While we should not expect busy citizens in a complex world to thereby become “omniscient,”⁷¹ maintaining an environment where people feel free to express themselves openly and without fear is essential to maintaining the limited-but-valuable form of accountability democracy is capable of sustaining. Meanwhile, liberalism curtails democracy's potential excesses by placing formidable barriers to the use of power, *including* that of the majority, so that the moods that sometimes sweep the public up in a given moment – such, I would argue, as the mood of pessimism that currently clouds our view of the democratic role and value of 21st century communications technologies – do not easily result in the silencing of valuable but currently unpopular voices or the overturning of public policies our better selves in better moments would have us endorse. For its part, democracy is the form of government most naturally suited to liberalism, since liberalism (as I will argue shortly) implies a strong optimism about the judgment of ordinary citizens, *in general and over time*. Democracy is also a match for liberalism in the sense that citizens that grow accustomed to life under liberal conditions tend to grow to cherish them highly, such that they become some of the primary goods the public is most keen to exert vigilance in order to protect.⁷²

I believe anyone who claims to defend liberal democracy on the definition offered and cares to an appreciable degree about public policy outcomes should agree with most of what is written above. At the very *minimum*, they need to agree that liberal democracy is *not itself likely to engender primary bads*, and certainly not to the degree that we need to abandon either the

70 One of the most famous empirically grounded defenses of “democracy,” Amartya Sen's argument that democracies do not allow their citizens to suffer famines, is actually a defense of *liberal* democracy. Sen's claim “we cannot find exceptions to this rule [that no democracy has ever experienced famine]” (Sen 1999, 5) may be too strong, as argued by Olivier Rubin among others (Rubin 2009), Sen seems to have identified a clear and strong correlation between democracy and famine prevention through the 20th century. Sen's claim is not, however, *exclusively* about democracy, at least not on my definition, but rather *liberal* democracy, since his claim is that “no substantial famine has ever occurred in any independent and democratic country *with a relatively free press*” (Sen 1999, 4-5, my emphasis).

71 This is the satirical word favored by Walter Lippmann, whose work examining the formidable problems with using public opinion to evaluate political performance is to my knowledge still unrivaled (Lippmann 1925, 21)

72 Perhaps the most dramatic recent demonstration of this phenomenon took place in Hong Kong, where millions of citizens accustomed to life under liberal conditions took to the streets for weeks to resist encroachments of the distinctly illiberal mainland Chinese government.

liberal or democratic part of liberal democracy to save ourselves from them. The tendency to resist primary bads is a consideration of paramount importance when it comes to choosing a political system. If a political system is likely to *engender* primary bads, its doing so constitutes strong a *prima facie* case against the desirability of that system. In this dissertation, I will argue that those I see myself as primarily opposing, the collection of scholars and public commentators I call Progressives, claim to support liberal democracy while consistently espousing and endorsing opinions that imply that liberal democracy is currently and should as a rule be expected to lead to primary bads. As we shall see, this leads them to support policies and embrace a conception of citizenship that are themselves illiberal, a move which I believe is both unjustified and quite troubling from a liberal democratic perspective.

Section Two: Doxastic Self-Rule and its Centrality to Liberal Democracy

Liberal Democrats' Strong Preference for Self-Rule

I have said that I will not supply a full-fledged theory of the kind of liberal democracy I support in this dissertation (though at this point I seem to be teetering on the edge of falling into one), but I *will* point out one key element of the version(s) of liberal democracy I support. This is the *strong preference for self-rule*. When I use the term “self-rule,” I just mean *making one's own decisions*, broadly conceived. I mean “self-rule” here very broadly, including both collective and individual self-rule. Liberal democracy values collective self-rule in the sense that it envisions the *demos* as exerting meaningful control over the public policies that govern it and are exercised in its name. This sort of self-rule has long been noted as integral to democracy, liberal or not. However, I believe liberal democracy's commitment to the desirability for *individuals* to exercise self-rule – for each individual citizen to come to their *own* conclusions and make their *own* decisions on the basis of their *own* judgment – is equally profound. My emphasis on the “their own” part does not mean liberal democracy envisions citizens as making decisions without being influenced by external influences, such as our upbringing or the opinions of others, the eradication of which from our decision-making processes, even if it were possible, would be undesirable. Instead, to say that liberal democracy strongly prefers for each individual citizen to make their *own* decisions is simply to say that liberal democracy desires for the decisive factor

that serves as the basis for citizens' decision is *each individual citizen's own judgment, exercised in the absence of coercion*. When I say liberal democracy has a *strong* preference for this kind of self-rule, I just mean that when the preference for self-rule conflicts with other of our preferences, it tends to win out. Liberals, for example, advocate for an extensive “private” sphere even though they acknowledge that harms to things we legitimately value regularly occur in that sphere. Likewise, many democrats advocate for universal suffrage even though they acknowledge that sub-optimal decisions, which can produce profoundly adverse effects on people's lives, are regularly reached by that decision-making method. Their willingness to suffer such results for the sake of honoring self-rule shows that their preference for self-rule is *strong*. Still, that does not mean it is *omnipotent*. In fact, it can't be, at least for the minimally consequentialist versions of liberal democracy I think are the only ones worth defending. The commitment to minimal consequentialism means that if citizens' exercise of self-rule is likely to result in political outcomes that do not meet minimal standards of acceptability whereas its abandonment would not, we would need to advocate for something other than self-rule. Liberal democrats believe this to be exceedingly seldom the case. In the overwhelming majority of situations, they view the abrogation of self-rule to produce worse outcomes than those the abrogation was intended to stave off. Because of this, liberal democrats tend to prioritize self-rule over other values into which it comes into conflict. This is why I say liberal democrats demonstrate a *strong preference for self-rule*.

Liberal Democracy's Modest Optimism about Lay Judgment and Strong Preference for Doxastic Self-Rule

There is no risk that a seasoned political theorist will be taken by surprise upon finding a self-identified liberal democrat expressing the desire to avoid the proliferation of coercion. While there is some disagreement about the precise definition of “coercion,”⁷³ most political theorists, and certainly most supporters of liberalism and/or democracy, have viewed coercion as something that can be justified by certain political necessities but whose influence over people's decisions we ought to either minimize or at least seek to prevent in key decision-making

⁷³ I am personally friendly toward the definition offered by Friedrich Hayek in *The Constitution of Liberty*, where he defines coercion as manipulating the alternatives available to a given decision-maker so that “the conduct that the coercer wants [them] to choose becomes for [them] the least painful one” (Hayek 2011 [1960], 200).

arenas, such as whom to vote for, which job to choose, which religion to follow, and where to live. We might call citizens' ability to make decisions in the absence of coercion in these and other ways *political* self-rule. In this dissertation, however, I am not primarily concerned with *political* self-rule. I take it for granted that most people, and certainly most of those sympathetic with liberal democracy as I have defined it, will already agree that political self-rule is desirable – that is, that it is desirable to avoid the proliferation of coercion in society, especially when it comes to the exercise of the rights mentioned above. The kind of self-rule I am primarily concerned with in this dissertation has been less emphasized in the writings of liberalism and democracy than the desire to avoid the proliferation of coercion, but I believe it to be no less central to most forms of liberalism, democracy, and liberal democracy. This is what I call *doxastic self-rule*. I define doxastic self-rule as the practice of adopting beliefs on the basis of *one's own judgment*, rather than through *routine deference to authority*. I will flesh out this definition and examine its implications more thoroughly in the coming chapters, but for now I just want to present a theoretical defense of the position that the desire that citizens exercise doxastic self-rule on important political matters *must* be central to any version of liberal democracy that is minimally consequentialist.

To argue that the sort of liberal democracy I have described is justified on *minimally consequentialist* grounds is to argue that both (a) constitutional enshrinement of core freedoms of communication, especially those of speech and press, and (b) the guarantee of universal, equal suffrage among all adults of sound mind are justified *at least partially on the grounds of their tendency to help society discover and embrace good (enough) political outcomes*, which, in turn, I have defined as manifesting a noticeable (if gradual) tendency to promote primary goods and a far greater than random tendency to avoid primary bads. I do not see how any minimal consequentialist in this sense could coherently embrace *both* (a) and (b) as integral features of a desirable political system without simultaneously endorsing the following two positions. First, they must believe that citizens should *act on the basis of their own judgment*, rather than *simply deferring to established and/or arbitrary authority*. At the very least, it is implicit in the very idea of minimally consequentialist liberal democracy that this be the case when it comes to in forming politically relevant beliefs and deciding how to vote, as I will argue presently. Second, minimally consequentialist liberal democrats must hold a certain *modest optimism* about lay

citizens' judgment, where “modest optimism” just means the belief that such judgment is sufficient to produce good enough political outcomes to justify going on with liberal democracy rather than jettisoning it for some alternative.

These conclusions are straightforwardly deducible from an examination of the decision-making conditions enshrined by liberal democracy in light of minimal consequentialism's requirement that public policy positions be justified at least in part by their tendency to produce good outcomes. The liberal component of liberal democracy seeks to guarantee to citizens a maximal ability to publish and access information, while the democracy component seeks to guarantee each of those citizens an equal say in fundamental political decisions. The minimal consequentialist could embrace either the liberal or democratic component *individually* without being either modestly optimistic about citizens' judgment or thinking they should act on its basis in making important political decisions. They could, for example, insist on freedoms of speech and press among the ignorant plebs as a harmless way of making them happier, so long as important political decisions are reserved to a select group of better decision-makers.⁷⁴ Or, as I will argue is far more common today, they could jettison their commitment to freedom of speech and press and retain the popular vote, reasoning that as long as the informational environment is cleansed of misinformation, fake news and propaganda, the public can be trusted to make good (enough) political decisions to keep going on with democracy.⁷⁵ Either way, it is possible for a minimal consequentialist to endorse *either* liberalism *or* democracy without entailing any necessary claims about the judgment of the lay public.

If the minimal consequentialist is to endorse *both* liberalism *and* democracy, however, they *must* both believe that (a) citizens should make important political decisions on the basis of their *own judgment* rather than automatically defer to pre-defined “authorities” and (b) citizens' judgment can be relied upon to produce good (enough) outcomes on at least the most important sorts of political decisions. The primary function of freedoms of speech and press would seem to be to guarantee *access to* as wide a possible range of information among the citizenry as well as to enable as broad a spectrum of citizens as possible the right to *disseminate* their opinions

⁷⁴ I take this to be the essence of the “epistocratic” regime defended by Jason Brennan (Brennan 2016).

⁷⁵ Again, this opinion is right in line with the work of Lippmann, who argued in *Liberty and the News* that without “protection against propaganda, without standards of evidence, without criteria of emphasis, the living substance of all popular decision is exposed to every prejudice and to infinite exploitation” (Lippmann 1920a, 37).

publicly. Meanwhile, the primary function of giving every able-minded adult citizen an equal vote would seem to be to give each of them *equal influence* on fundamental political decision-making processes. If we did not want citizens to rely on their own judgment to decide what to believe, and instead wished for them to defer to some set of authorities, why would we risk guaranteeing them *access to* as wide as possible a range of information and allowing them to *freely disseminate* their ideas to the rest of the public? And if we did not want them to rely on their own judgment to decide how to vote, and instead wished for them to defer to the opinion of some set of authorities on the matter, why wouldn't we just leave that decision up to the very authorities to whom that deference is desired? To maintain liberalism's insistence that lay citizens ought to have access to as wide a range of information as possible and democracy's guarantee to those same lay citizens an equal opportunity to influence fundamental political decisions in the way characteristic of democrats seems entirely baffling, even reckless, unless one wishes for those citizens *not* to merely defer to some set of pre-defined authorities but instead to adopt beliefs and make decisions on important political matters on the basis of their own judgment. In other words, *minimally consequentialist liberal democracy necessarily implies a preference for doxastic self-rule*. Moreover, the force with which liberal democrats defend freedoms of communication and universal, equal opportunity suffrage implies that this preference must be quite *strong*. In other words, *liberal democracy necessarily implies a strong preference for doxastic self-rule*.

Once this is clear, it is an easy step to conclude that minimally consequentialist defenses of liberal democracy imply at least a modest optimism about the judgment of lay citizens. This can be seen by considering the consequences a deep pessimist about the judgment of the lay public must expect from the maintenance of freedoms of communication and universal, equal opportunity to participate in important political decision-making processes. For anyone deeply pessimistic about the public's judgment on important political matters, liberal democracy's combination of free and open communication and universal, equal opportunity to take part in fundamental political decisions would seem to be at best blind submission to *fortuna* and at worst an active invitation to destructive public policies, neither of which should be acceptable to the minimal consequentialist. Thus, in order to defend liberal democracy as (at least) a minimal consequentialist implies (a) a belief that citizens ought to rely on *their own judgment* when it

comes to adopting beliefs and making decisions on important political matters and (b) at least a certain *modest optimism* about the judgment of lay citizens – specifically, that their judgment can be relied upon to secure minimally acceptable political outcomes.

Legitimacy, Justification and Judgment

From what I have written so far, it should be clear that a certain modest optimism about lay citizens' judgment and a strong preference for doxastic self-rule are central to any minimally consequentialist defense of liberal democracy. I have argued that any model of liberal democracy worth defending is minimally consequentialist, so it follows from just this argument that I believe a certain modest optimism about the judgment of lay citizens and a strong preference for doxastic self-rule is central to any liberal democracy worth defending. But there is another element, beyond minimal consequentialism, that I and many others believe ought to be part of any model of liberal democracy if it is to be worth defending, which further cements the place of this modest optimism about lay judgment and the strong preference for doxastic self-rule at the heart of those models. This is the requirement that *in order for demands of deference to authority to count as legitimate, those demands must be justified to those whose deference is demanded through reason*. I call this the *justified deference principle*.

The justified deference principle is a necessary but not sufficient condition. A demand of deference is not *automatically* legitimate *just because* the demander accompanies that demand with a reason or reasons, but *no* demand of deference to authority can be legitimate if it is *not* justified by reason(s) given to those whose deference is demanded. Stated as such, I believe the justified deference principle is an integral part of any liberal democracy worth defending. Part of the appeal of both liberalism's valuation of *universal, equal* rights and democracy's valuation of *universal, equal* suffrage is that both resonate with the assumption of fundamental moral equality shared by many of us in the contemporary age. Most of us believe that people ought by default to be treated as if they are of equal worth and shown equal dignity. In large part, we affirm this default equality of worth and desert of dignity by allowing people to make their own choices – choices about what religion to practice (or not), which beliefs to subscribe to, which causes to pour their energy into and, of course, which public policies and political candidates to vote for. Justifying demands of citizens' deference to authorities affirms citizens' dignity in the same way.

It demonstrates that their deference is appropriate not because they inhabit some subordinate position or belong to some subaltern caste, but because deference to the given authority *makes sense* the situation at hand according to reasons those citizens are capable of understanding and acknowledging. The age-old slogan “no taxation without representation” is an example of our embrace of the justified deference principle. It resonates with us because it affirms that those whose payment of taxes is demanded ought to know and approve of how those taxes are being used. Implicit in this is the idea that those taxpayers are *capable of judging for themselves* whether their tax money is being used well enough to *justify* the demand that they pay it as a matter of law, or at least whether the representative they select to defend their interests to the tax-collecting authority is likely to be a good defender of those interests.

In recent years, liberal political theorists have written a great deal about the role of reason in justifying demands of deference to authority. This is the central issue at stake in the literature on “public reason” that flourished among liberal theorists, particularly those identifying as “political liberals,” during the last few decades. The unifying idea for theorists who subscribe to the ideal of public reason is that the reasons offered to citizens as justifications for demands of deference ought to be *public*. What it means for reasons to be “public” in this way is a matter of debate, but the general idea is that, as Simone Chambers puts it, the reasons given to citizens to justify demands of their deference ought to be “intelligible, accessible, acceptable, or sharable.”⁷⁶ Not only, then, do public reason endorsers believe that demands of deference to authority need to be *justified*, but that only a *certain kind* of justification – justification that meets the standard of *public reason* – can actually make those demands legitimate. Most public reason theories are based on what might be called the idea of *hypothetical consent*. This is the idea, in Chambers' words, that “[b]ecause we cannot as an empirical matter canvas every single person and also because actual responses of empirical persons can be unreasonable (the traditional problems with consent), we ought to seek to fold others [*sic*] people's possible consent...into our very reasoning.”⁷⁷ In other words, because we can't ask every single citizen whether or not a given justification makes sense to them and because citizens can be unreasonable even if we could, most public reason theorists write as if it is the task of political theorists and public servants to

76 Chambers (2010), 894

77 Ibid., 898

try to determine for themselves what ought to count as publicly reasonable and use that as a basis for demands of deference. This is indicated by the common “-able” and “-ible” endings on all the adjectives chosen by Chambers to illustrate the attributes of reasons liberal theorists have emphasized in order to count as “public reasons.” It is also evident in the way Thomas Nagel describes the public reason theory of John Rawls’ as basing institutions’ legitimacy “on their conformity to principles which it *would be* reasonable for disparate individuals to agree upon”⁷⁸ and T. M. Scanlon’s view as demanding that social rules be only those which no one with pure motives “*could* reasonably reject.”⁷⁹

I agree with public reason theorists’ insistence that demands of deference to authority must be justified by reasons addressed to those whose deference is demanded. I also think the concept of hypothetical consent often emphasized by public reason theorists encourages habits of perspective-taking that are valuable in all aspects of social life, especially in diverse societies like most contemporary democracies. However, one downside to public reason theorists’ heavy emphasis on hypothetical consent is that it cuts the *public* largely out of the process of “public” justification. Too often, public reason theorists write as if the only or primary value of the concept of public justification is to serve as the basis for a sort of thought-experiment for theorists and/or policy makers, whose job is to imagine *for themselves* what sorts of reasons *ought* to be acceptable to *everyone else*. The goal seems to be for theorists and policy makers, by engaging in this type of thinking in a sincere and rigorous way, to deduce what sorts of reasons *ought* to be acceptable to all “reasonable” citizens, thereby allowing themselves to make demands of citizens in full confidence of their legitimacy (since they meet the standard of public reason). At no point in this process is it implied that the *judgment of lay citizens themselves* might have anything to do with whether the demands made of them ought to count as legitimate. In fact, this approach seems very often to be infused with a certain doubt about the fitness of ordinary citizens to make such judgments, as evidenced both by liberal theorists’ tendency to quickly dismiss the viability *empirical* consent as a standard of justification in favor of the much more emphasized standard of *hypothetical* consent, as well as their preoccupation with how to incorporate “unreasonable” individuals and publics into their models of public justification.⁸⁰

78 Nagel (1987), 220

79 Ibid., 221

80 E.g. Van Schoelandt (2015)

While the concept of hypothetical consent can be useful insofar as it encourages perspective-taking, I believe encouraging political theorists and/or policy makers to approach political issues as if they are capable of determining *by themselves* what the “true” demands of public reason are is a mistake. It is a mistake both because, as Charles Sanders Peirce wrote in a critique of Cartesianism long ago, “to make single individuals absolute judges of truth is most pernicious,”⁸¹ as well as because it limits our view of the value of public justification by focusing on the lesser of its functions. The notion of public justification ought *not*, in my opinion, be solely or primarily a means of encouraging political theorists and policy makers to try to engage in the right sorts of thought-experiments. Instead, it ought to serve as a reminder that in all instances in which citizens' deference is demanded the citizens ought to be given *reasons* for that demand *so that they can judge for themselves whether the demand itself is legitimate*. On this view, the “public” part of the term “public reason” is better conceived as a synonym for “publicized” than as a synonym for “generally applicable,” and the term “public reason” ought to be a reminder of the central role played by the *public's* judgment in determining whether demands for its deference are legitimate rather than implying that the judgment at the heart of determining their legitimacy is the judgment of the *theorist*.

I believe this notion of the related terms “public justification” and “public reason” capture a core component of what is unique and valuable about the form of liberal democracy I view as worth defending. It embodies the combination of respect for the dignity of lay citizens and modest optimism about the judgment of lay citizens I have argued to be both *implied by* support for liberal democracy and to be *supported by* its historical performance. I also believe the justified deference principle and the standard of public reason are appropriate normative standards for *both* political deference, in terms of submission to the law, as well as *epistemic* deference, which I define as the acceptance of or agreement to act on the basis of the claims of some given authority. No matter whether we are demanding that citizens follow an order or adopt a belief, our demand that they do so is not *justified* unless it is accompanied by *reasons addressed to those citizens whose deference is demanded*. The reason this is so important is because it facilitates those citizens' ability to *judge for themselves* whether our demands meet the appropriate normative standards to merit their deference. This provision of reasons invites liberal

81 Peirce (2006 [1868]), 71

citizens to exercise doxastic self-rule as a crucial element of the process of justification. Conceived this way, the notions of public justification and public reason that have occupied liberal theorists in the past few decades also affirm a central role for the exercise of doxastic self-rule by lay citizens in a well-functioning liberal democracy.

The Incompatibility of The Progressive View with Minimally Consequentialist Liberal Democracy

Having now explained and defended my views about liberal democracy, I will move on in the following chapters to explain why I think the Progressive view is both incompatible with many Progressives' avowed support for liberal democracy and why we in the 21st century should embrace public policies and notions of civic epistemology that *are* consistent with liberal democracy, rather than those subscribed to by Progressives. I will close this chapter by clearly laying out my main objections to the Progressive view in terms of its incompatibility of liberal democracy, which will be elaborated in greater depth over the course of those chapters. Progressives run afoul of the sort of liberal democracy I have described in this chapter by:

1. Consistently expressing a deep pessimism about the reliability of the public's judgment on matters of political importance that is strikingly at odds with the modest optimism I have argued to be implicit in convincing defenses of liberal democracy
2. Discouraging and disparaging the exercise of doxastic self-rule by lay citizens, indeed often arguing it to be *itself* a threat to liberalism and/or democracy
3. Espousing views that clearly imply that liberal democracy is very likely to facilitate primary bads in the form of the large-scale endorsement of false and/or hateful beliefs, rule-by-demagogue, and easy manipulation of the public by bad actors in pursuit of destructive political ends
4. Championing approaches to public policy issues related to 21st century communications technologies that violate the core tenets of liberalism both in letter and spirit, such as endorsing governmental and/or corporate censorship of citizens' information online
5. Demanding citizens' deference to institutional authorities without any attempt at accompanying those demands with reasons addressed to the citizens whose deference is demanded, thus violating the standard of public justification

In the final chapter of this dissertation, I will argue that, compared to to Progressives, the sort of citizenship practiced by citizens Progressives disparage as being afflicted by the “post-truth malaise” is considerably more compatible with the ideals of liberal democracy on all these fronts. These commentators tend to:

1. View the public – i.e. themselves – as capable of bearing the “burdens of judgment” minimally consequentialist liberal democracy necessarily places upon them
2. Actively encourage citizens to exercise doxastic self-rule by insisting that deference to *any* authority be justified in the light of each citizen's own judgment, rather than expected as a matter of course when it comes to certain authorities pre-defined as “reliable,” “trustworthy,” or “vital to democracy”
3. View the exercise of doxastic self-rule by lay citizens to be a *defense against* primary bads, prominently the proliferation of corruption and illegitimate forms of authority
4. Strongly oppose flagrantly anti-liberal public policies, especially censorship of communication online
5. Insist that demands of deference to authority be consistently and without exception accompanied by reasons directed to the citizens whose deference is demanded

In the following chapters, I will explain how Progressives manifest these tendencies on a wide range of issues, all of which are centered in one way or another on the impact of 21st century communications technologies on democracy. As we will see, Progressives are quite pessimistic about the impact of those technologies, but I believe this pessimism is unjustified. One of the main reasons is that the Progressive idea of what it means for democracy to be “doing well” is poorly specified and built on Progressives framing their expectations about and interpretations of the democratic impact of 21st century communications technologies around an excessively pessimistic view of lay citizens. To this argument I will now turn.

Chapter Two: The Epistemic Well-Being of Democracy

“The epistemic well-being of democracy” is a term that has never gained currency either among scholars or the public at large. While this might speak well of both groups' aesthetic sensibilities, it might also be a missed opportunity, for, although it is seldom discussed in such terms, a number of high-stakes political discussions on contemporary issues revolve around worries about democracy's continuing epistemic well-being. The advent of new technologies like the internet and social media and their corresponding impact on the ways individuals and institutions disseminate, filter and process information have given rise to a number of such worries. Will individual citizens set upon by the beleaguering array of informational sources made available in the Information Age struggle to tell truth from fiction?⁸² Will the ability to select one's own group of peers result in ignorance of or unwillingness to consider the views of others?⁸³ Will the decline of the legacy news industry in the wake of vast shifts in our communications technologies make it harder to hold power to account?⁸⁴

Such questions have caught the imagination of scholars from an array of disciplines as well as lay commentators writing for in an array of publications, and all have to do with apprehensions about the epistemic well-being of democracy. It is difficult, when one has read enough of this commentary, to miss the mood of pessimism that tends to dominate. Where a decade ago the internet and social media were viewed as harbingers of a new and glorious democratic era, today they tend to be discussed in bitter tones, with many worry aloud that these technologies, once heralded as democracy's 21st century allies, may in fact be its enemies. Much of this worry stems from the belief that in order for democracy to be doing well, the informational environment needs to be functioning properly. That is to say, people need access to reliable information if they are to recognize and agree about the problems that exist within our

82 Versions of this worry are implicit in the ascendancy of “misinformation” and “fake news” as topics of concern. Examples include Kovach and Rosenstiel (2011), Minozzi (2011), Lewandowsky (2012), “Global Risks” (2013), Deb, Donahue, and Glaisyer (2017), and Schiffrin (2017).

83 This worry has been most persistently raised by Cass Sunstein. See, e.g., Sunstein (2017).

84 See, e.g., Alterman (2012), Derakhshan (2019)

society and endorse good solutions to them. The advent of the internet and social media have, it is often supposed, made identifying reliable information more difficult. Worse, they have made finding *misinformation* easier.⁸⁵ Because of these and other phenomena related to the rise of the internet and social media, it would seem that contemporary democracy is in the eyes of many commentators faring pretty badly in epistemic terms, and that a major driver of this poor performance is changing communications technology. A widespread sentiment seems to be shared that the emergence of new communications technologies, especially social media, represents a threat to the epistemic well-being of democracy, to the extent that numerous commentators even suggest that democracy *itself* is under threat.⁸⁶

In this chapter, I will oppose this pessimism. I will argue that the dismal mood that currently dominates discussions about 21st century communications technologies' effect on the epistemic well-being of democracy is based on a rough and incomplete theory of what it might mean for democracy to be epistemically “doing well.” This incomplete theoretical picture has led not only to premature pessimism about the epistemic well-being of democracy and the role played by 21st century communications technologies in fostering it, but has also led scholars to pursue a research agenda that is not designed to adequately answer the question of whether new technologies like the internet and social media constitute threats to the epistemic well-being of democracy. In fact, I will argue, the current empirical research agenda is so skewed that it is likely to lead to the conclusion that the epistemic well-being of democracy is diminishing even if it is actually *improving*. In part, this is because the commentators espousing these views appear to *presuppose* the sort of *deeply pessimistic* view about the judgmental capacities of lay citizens I argued in the previous chapter to be incompatible with any convincing defense of liberal democracy, and to interpret the likely effect of 21st century communications technologies on democracy in that pessimistic light.

The essay will proceed as follows. In Section One, I will show that, although commentators rarely, if ever, explicitly explain what it would mean for democracy to be doing well epistemically, many of those who are most worried about the epistemic well-being of democracy are animated by common ideals, whose combination I will call the *Progressive Ideal*.

85 Kuklinski et al (2000), Lewandowsky (2012)

86 Deb, Donahue and Glaisyer (2017), Persily (2017), Schiffrin (2017), Levitsky and Ziblatt (2018)

Though I have some objections to the Progressive Ideal, I will save them for later chapters. Here, I will restrict myself to the attempt to show how, even if one accepts the Progressive Ideal, the pessimism that often dominates lay and scholarly discussion about 21st century communications technology's impact on the epistemic well-being of democracy is not clearly justified by the evidence currently on offer. In order to do this, I will develop a list of conditions that must be true of a democratic society's informational environment in order for it to foster the epistemic well-being of democracy according to the Progressive Ideal. Once these conditions are stated explicitly, I will move on to Section Two, in which I identify and explain four criteria research must meet if it is to give us insight into how well democracy fulfills these conditions. In Section Three, I will show that at least one of these criteria, and often several simultaneously, are neglected by commentators who view the internet and social media as threats to the epistemic well-being of democracy, in large part because such commentators consistently frame their interpretations of the impact of 21st century communications technologies around a deeply pessimistic view of the judgmental capacities of lay citizens that is far from clearly justified. In each case, we will see how, if we decline to start off with this deep pessimism and take into account *all* of the four criteria listed above, the Progressive conclusion that the epistemic well-being of democracy is threatened by technology-driven changes to the 21st century's communications environment into serious doubt. In fact, in many cases – some considered in this chapter and some to be considered in greater depth in chapters to come – a more thorough analysis reveals that we have just as much reason to believe the exact opposite, that is, that the effect of the internet and/or social media has actually been to *improve* the epistemic well-being of democracy along the same dimension for which they have been criticized for harming it.

In this chapter's conclusion, I will argue that because so much of current commentary on and research about the epistemic well-being of democracy is based on an incomplete idea of what it might mean for democracy to be doing well epistemically, and because the very same evidence cited by Progressives as evidence of democracy's epistemic decline may often be evidence of just the opposite, Progressives' pessimism about the impact of 21st century technologies on democracy, influential and pervasive though it may be in the present day, is at best premature and and worst fundamentally misguided. More importantly for my purposes, it undermines Progressives' case for adopting the deeply pessimistic view of the judgmental

capacities of lay citizens that both contributes to and is built upon their interpretation of 21st century communications technologies' impact on democracy. Not only does this mean that commentators should resist the urge to treat 21st century communications technologies, and the lay citizens who use them, with the deeply pessimistic air that is currently in vogue, but it also means that the considerable amount of pressure being brought to bear by concerned commentators for authorities both private and public to exert more control over tech and social media companies may be misguided, perhaps even severely so. Given the dangers entailed in giving the public's blessing to political or extra-political curation of the information to be made available to the public, we should resist such policies unless and until clearer evidence, interpreted in the light of a clearer and more well-rounded view of what the epistemic well-being of democracy entails, has been gained.

Section One: The Legacy of Progressivism and Democracy's Perceived Contemporary Epistemic Problems

Contemporary Worries about the Epistemic Well-Being of Democracy

Misinformation⁸⁷ has been something of a hot topic among psychologists, political scientists and communications scholars for the past two decades. Its ascendancy owes in part to the rapid changes that have taken place in communications technology in that period. The swift rise and diffusion of the internet, smartphones and social media has dramatically decreased the costs of both *publication of* and *access to* information. These decreased costs have allowed an enormous variety of publicly available informational sources to emerge and has allowed citizens to more easily and cheaply access those sources. At the same time, these developments have undercut the “gatekeepers” – radio and TV station producers, newspaper and magazine editors, book publishers and others – that once played a crucial role in curating the information to which lay citizens had access. This explosion of accessibility and dramatic reduction in both the costs of publication and the gatekeeping role played by traditional publishers have given rise

⁸⁷ Some scholars (e.g. Karlova and Fisher 2013) make a point of distinguishing between “misinformation” and “disinformation.” The former is a catch-all category including inaccurate information that is *either* deliberately *or* unwittingly transmitted, while the latter refers only to inaccurate information that is knowingly and deliberately transmitted. I mean to include both these kinds of inaccurate information when I use the term “misinformation.”

to worries that the average quality of the information accessed by citizens today might decline as compared with those of yesteryear, with negative consequences for the decision-making ability of democratic citizens.⁸⁸ These sorts of worries have become more widespread and visible in the last several years, particularly since the 2016 election cycle, which saw the term “fake news” explode into public consciousness. Though its meaning has since become politicized,⁸⁹ when it originally broke into public consciousness during and after the 2016 American presidential election cycle, the term “fake news” referred to false stories done up in the style of mainstream news. For many commentators it epitomized the difficulty faced by citizens attempting to discern truth from falsehood, particularly in the arena of politics, in the 21st century's rapidly changing informational environment.⁹⁰ These worries about misinformation, fake news, and the overall quality of the information whose spread is facilitated by 21st century technologies are obviously epistemic in nature. They are also usually *democratic*. Both in the academy and the mainstream, many of those who express worry about the impact of 21st century communications technology worry because they think it will be difficult for citizens who are served by the open and informationally dense epistemic environment created by the internet to sort through the chaos in a way that allows them to form accurate beliefs in the way necessary for them to vote appropriately and to recognize and agree upon the solutions to important societal problems. In other words, they are worried about the impact of the internet and social media because they see these phenomena as posing a threat to what I will call *the epistemic well-being of democracy*.

Yet, even though this kind of concern for the epistemic well-being of democracy is the obvious normative driver for a great bulk of the burgeoning scholarly and public commentary on

88 This sort of concern about fake news is widespread and has been given voice by scholars in multiple fields, public policy advisors and analysts in multiple countries, and editorialists in numerous mainstream publications. What follows is a non-exhaustive sample intended to illustrate the range of commentators who have explicitly voiced worries that the decreased costs of publication and decline of traditional “gatekeepers” have decreased the average quality of information accessed by democratic citizens. In the academy, see Goldman (2010) for an example in philosophy, Garrett (2017) for an example in communications, Cannon (2001) for an example in journalism, and Schiffrin (2017) for an example in political science. Among policy advisors and analysts, see the British House of Commons Special Report on Disinformation and Fake News (2018), UNESCO's report on Journalism, 'Fake News' & Disinformation (2018) and Deb, Donahue and Glaisyer (2017). And for just two of the voluminous examples in the mainstream punditry sphere, see Jackson (2018) and Alary (2017).

89 Elsewhere (Whitsett 2019), I have argued that “fake news” is used in two largely contradictory ways. In its *literal* sense, the term “fake news” is used to describe made-up articles published online in a style that mimics that employed by professional journalists, while in its *rhetorical* sense the term is used as a shorthand dismissal of the reliability of a given source or report. The sense I refer to in this passage is the literal one.

90 For the purposes of this chapter, uses of the term “fake news” refer to what I have in a previous chapter dubbed the “original sense” of fake news unless otherwise specified.

misinformation, fake news, and social media's overall impact on politics, an explanation of what exactly it would *mean* for democracy to be doing well epistemically is very seldom, if ever, articulated in any detail. As we will see, I think the absence of such an explanation contributes to some problems in that literature; however, I do *not* think its absence means there is no common normative framework informing this literature at all. On the contrary, I believe that many of those who are most concerned about misinformation and fake news are animated by common and reasonably coherent ideas about what it means for democracy to be doing well epistemically. The problems that arise from commentators neglecting to explicitly state what it would mean for democracy to be doing well epistemically are not due to the absence or incoherence of the normative ideals that motivate their concern, but are instead due to other problems that commonly arise from failing to be explicit about one's theoretical orientation, such as encouraging one to draw excessively broad conclusions from a narrow range of evidence or to fail to see the way a given fact could be evidence for a given conclusion *in one sense* yet at the same time be evidence for an opposite conclusion in another. In order to show what is wrong with the normative framework that guides these commentators' thinking, I must first articulate what I take that shared normative framework to be.

The Legacy of Progressivism

By far the most common normative tie that binds contemporary commentators who worry about the epistemic well-being of democracy is a general endorsement of the truism “democracy depends on a well-informed citizenry.” Not only is that truism invoked with some frequency in both academic and non-academic work on misinformation,⁹¹ but it has served as a bedrock normative assumption in Americans' folk ideas about how democracy ought to work for many generations, ever since the ideal of the “informed citizen” became popularized by the original Progressives in the late 19th and early 20th centuries. The important points of the Progressive legacy for the purposes of this essay are (1) its emphasis on individual citizens having access to

91 The mission statement on TruthOrFiction.com's “About” page, for example, begins with the statement that “Democracy cannot exist without an educated and informed populace” (TruthOrFiction 2019). This sentiment aligns perfectly with the assertion of Kuklinski and colleagues that for the last half-century at least “the normative thrust in public opinion research has been unwaveringly: citizens should be factually informed” (Kuklinski et al 2000, 790).

and voting on the basis of *reliable information*⁹² about the state of the world and the policy platforms advocated by political candidates, rather than on the basis of partisan identity and/or venal self-interest, as had been the case for much of the 19th century; and (2) the conviction that reliable information is most often disseminated by *professional* and/or *expert*⁹³ sources, primarily because these kinds of information-seekers obtain their knowledge via some especially sound or “objective”⁹⁴ truth-seeking method(s).⁹⁵ Together, I will call these two tenets the *Progressive View*. From these a *Progressive Ideal of Democratic Epistemic Well-Being* can be derived:

Progressive View: Democracy's epistemic well-being depends crucially on (1) citizens having access to and acting on the basis of *reliable information*, where (2) reliable information is most often disseminated by the relevant *professional* and/or *expert* source(s).

Progressive Ideal of Democratic Epistemic Well-Being: A state of affairs in which citizens have access to and generally act on the basis of reliable information, which importantly includes a propensity to consult and believe expert and/or professional knowledge-disseminating institutions.

92 I will occasionally use terms like “fact” or “truth” in this essay, but unless explicitly stated otherwise these words are intended in a way that takes into account the fallible and contingent nature of all human knowledge, which I take to be the dominant view of the nature of human knowledge among all commentators. The term “reliable” is often used, not as an endorsement of the epistemic school of reliabilism, but simply to make clear that I take most everyone to believe that (a) all our knowledge is fallible yet that (b) we can be more justifiably confident about some things than others.

93 By no means do I wish to imply that “expert” and “professional” are synonyms. Instead, I mean to voice my agreement with epistemologist David Coady (Coady 2012, ch. 6), who points out that both professionals and experts are often treated as playing an equivalent, and vital, role in curating the information it is important for citizens to place their trust in if democracy is to function properly epistemically.

94 I want to stress that I mean “objective” here to refer to the *methodological approach* used by knowledge-seekers like scientists, journalists and a variety of other professionals in the contemporary era. Without wanting to get bogged down in a notoriously complicated topic, I'll just say that the kinds of methods I'm talking about are generally intended to do away with or in some other way compensate for common biases that distort humans' pursuit of knowledge, like fidelity to an ideology or the acceptance of unacknowledged and limiting assumptions. The adoption of these methods is often viewed as a way of securing knowledge of a more dependable sort than that obtained via less rigorous methods.

95 These characterizations of Progressives and their ideals, as well as the advent of the modern notion of “objectivity,” are quite conventional. For a good summary of the Progressive movement and its impact on American ideals of the “good citizen,” specifically, see Schudson (1998), ch. 4: “The Second Transformation of American Citizenship.” For a historical account of the rise to common parlance of the notion of objectivity, see Daston & Galison (2007).

Communications historian Michael Schudson argues that American citizens have inherited our default notions of “good citizenship,” at least when it comes to epistemic matters, from the Progressives, calling our default view the “century-old Progressive Era model of the informed citizen.”⁹⁶ Whether or not this is true for citizens at large (I suspect it is, in the United States at the very least), the Progressive View continues to be the default ideal of the epistemically good citizen among most journalists as well as many social scientists. More to the point, it is clearly the normative ideal endorsed by many of the social scientists, editorialists and politicians who are worried about misinformation and fake news, who I have – not coincidentally – chosen to call “Progressives” in this dissertation. Sometimes, acceptance of the Progressive Ideal is expressed more or less explicitly, as when in a recent article Stephan Lewandowsky and colleagues agree that it is a “truism” that “a functioning democracy relies on an educated and well-informed public”⁹⁷ and emphasize the societal importance of citizens being willing to trust professional sources of knowledge, such as “official information”⁹⁸ and that produced by “experts,”⁹⁹ which they take to be based on “facts and objective evidence.”¹⁰⁰ More often, however, either one or both parts of the Progressive View are simply taken for granted. For instance, when James Kuklinski and colleagues report that “the normative thrust in public opinion research has been unwavering: citizens should be factually informed,”¹⁰¹ making part (1) of the Progressive Ideal they and other public opinion scholars endorse explicit, they do not immediately explain what *kind* of “factual informedness” citizens need to possess in order for democracy to work properly. But what they *mean* by “factual informedness” can be inferred from their method for obtaining the “factual items” they use to form their tests of civic informedness. To gather such facts, Kuklinski and colleagues report having consulted with “welfare experts” and used their claims about welfare policy as the metric against which participants' knowledge was compared.¹⁰² The implication is clear. Factual knowledge is the kind produced by *experts in the relevant field*, and it is against this standard that citizens' level of “informedness” is to be assessed. Just like Lewandowsky and colleagues, then, the normative

96 Schudson (1998), 75

97 Lewandowsky et al (2017), 354

98 Ibid., 355

99 Ibid., 354

100 Ibid., 361

101 Kuklinski et al (2000), 790

102 Ibid., 795

assumptions of Kuklinski and colleagues fall in line with what I have called the Progressive Ideal.

Like these scholars, many commentators who are especially worried about the ways in which technological developments facilitate the flow of misinformation and fake news seem also to embrace one or both of what I have called tenets (1) and (2) of the Progressive View.¹⁰³ Nor are these worries and the normative ideals that motivate them reserved to the musty confines of the Ivory Tower. On the contrary, they can be spotted in abundance in the speeches of politicians, the reports drafted by political advisory groups, and the writings of editorialists. UNESCO's Handbook on Journalism, Fake News and Disinformation, for example, claims that "Knowledge that [is] news – produced by transparent actors and which is verifiable – is essential for democracy,"¹⁰⁴ and elsewhere names professional journalists as "communicators who work in the service of truth" whose work is "needed as an alternative, and antidote, to the contamination of the information environment."¹⁰⁵ Here again we can see the belief that reliable information is essential to democracy, as well an affirmation of the importance to democracy of the knowledge produced by a given set of dedicated professionals. Both in the academy and the mainstream, concern about 21st century technology's impact on the epistemic well-being of democracy goes hand-in-hand with, and is often in some way justified by reference (implicit or explicit) to, the two tenets of the Progressive Ideal.

Though I have some objections to this part of the Progressive legacy, I will not air them here. Rather than to critique, my main goal in this section has simply been to show that acceptance of tenets (1) and (2) of the Progressive View, and a corresponding general endorsement of the Progressive Ideal, are common among commentators who are worried about the epistemic well-being of democracy. For the rest of this essay, I will for the most part accept these tenets without reservation, saving any criticism of these and related ideals subscribed to by Progressives for other chapters. Instead, I will use the rest of this chapter to show that *even if* one accepts these tenets of Progressivism, it is not clear that 21st century communications

¹⁰³Legal scholar Nathan Persily, for instance, in an article whose title asks whether democracy can "survive" the internet, warns that the spread of fake news during the 2016 election campaign may have revealed the "vulnerabilities of democracy in the Internet age, especially when it comes to the integrity of the information voters will access" (Persily 2017, 67-9). See also Roth (2012) and Lewandowsky et al (2012, 107).

¹⁰⁴UNESCO (2018), 13

¹⁰⁵*Ibid.*, 8

technologies pose a threat to the epistemic well-being of democracy. In fact, the very same evidence emphasized by Progressives may just as well lead us to the opposite conclusion – that 21st century communications technologies are being employed by citizens in a way that is *beneficial* to the epistemic well-being of democracy. Because of this, we have good reasons to be wary of Progressives' negativity toward 21st century communications technologies and the deeply pessimistic view of lay citizens that negativity contributes to and is partly built upon, and should accordingly be hesitant to embrace Progressives' favored approach to both responsible citizenship and public policy in the digital age.

Section Two: What it Means for Democracy to be Doing Well Epistemically

Practical Implications

Even if we agree with both the Progressive View and the Progressive Ideal, neither is going to get us very far in the real world. That's because both tenets are just abstract axioms. Like mathematical axioms, they neatly summarize principles that we think apply somehow to the real world but, also like mathematical axioms, *by themselves* they bear little relevance to our lives. If we leave the tenets of Progressivism in this abstract form, we are left without (a) any *practical guidance* for how to live our lives and construct our public policies and (b) any means of *assessing the desirability and/or viability* of the Progressive Ideal itself.

Take, for instance, tenet (2) of the Progressive Ideal, the view that people in general ought to be able to access and willing to consult and defer to the claims of experts and professionals in forming their beliefs. Though I imagine plenty of people would agree with this in the abstract, without knowing *which* individuals and institutions *count* as experts/professionals, tenet (2) does us little practical good no matter how vociferously we endorse it. Or take tenet (1), that people ought to have access to and be willing to believe reliable information. Even fewer people would disagree with this than tenet (2). But, again, what *counts* as reliable knowledge is left entirely undefined. Moreover, nobody can be *fully* informed about *everything* relevant to politics, yet the Progressive Ideal tells us nothing about how citizens ought to prioritize their knowledge-seeking activity. Nor are these trivial omissions. On the contrary, the questions of what *counts* as “reliable knowledge,” *how much* of which *particular things* democratic citizens need to know

about politics, and how we're supposed to identify the *real* experts are each at the center of long-running and seemingly interminable debates among political scientists and philosophers and, as we shall see, bear important implications both for our views about what counts as responsible citizenship and for the sorts of public policies we embrace with regard to the management (or non-management) of our informational environment.¹⁰⁶

Earlier, I noted that most pessimists about 21st century technologies' impact on the epistemic well-being of democracy do not even explicitly articulate the *ideal* that drives their concern. Now we can see that this omission is just the tip of the under-theorized iceberg. Even if they *had* articulated the Progressive Ideal, if they did so without supplementing that ideal with some kind of concrete guidance about *how much* of *which types* of political knowledge citizens need to have or *which people/institutions ought to count* as professionals/experts, it would still be nearly impossible for us to diagnose the epistemic well-being of democracy according to either tenet (1) or (2). In other words, the Progressive Ideal cannot *by itself* be used to assess the epistemic well-being of democracy. To do that, we need to figure out ways to use the ideas implicit in the Progressive Ideal to assess the epistemic well-being of democracy via concrete, observable phenomena. In other words, we need to develop *standards* for empirically assessing the epistemic well-being of democracy according to the Progressive Ideal.

Progressive Standards for Assessing the Epistemic Well-Being of Democracy

Given the apparently intractable state of disagreement that plagues political scientists and philosophers about how much of which kinds of information citizens ought to have and who counts as the “real” experts, developing the kind of standards mentioned above might seem

¹⁰⁶Epistemology may be fairly described as the organized Western practice of arguing over the question of what counts as reliable knowledge. The question of what *kind* of knowledge citizens need to possess has embroiled political scientists in arguments for decades. As it stands, Doris Graber's statement that “there is no uniform, widely accepted answer to that question” remains as accurate today as when it was written nearly two decades ago (Graber 2003, 149). The main players in the debate have been those who wish the public to be as informed as possible on substantive issues, such as Delli Carpini and Keeter (1996) and those who deny that such knowledge is especially important due to the efficiency with which citizens can use heuristics to make decisions “as if” they had such knowledge (e.g. Popkin 1994). As for the question of how lay citizens are to figure out who the real/reliable/trustworthy experts are, it has proven similarly vexing for political theorists and philosophers, with many skeptical of whether it is even possible for them to do so reliably at all. See Turner (2003) for an excellent analysis of just how problematic the problem of expertise is for liberal democrats in particular. For a recent article arguing that in political matters, at least, telling who are the real (i.e. correct) experts is nearly impossible, see Viehoff (2016).

hopeless. Fortunately, however, I think we can avoid getting caught up in these age-old debates if we make no attempt to construct an account that is *generally* acceptable, and instead restrain ourselves to a couple of beliefs that are widely shared by Progressives. One of these is a belief in the importance to democracy of the legacy news media and the professional journalists employed by them. Recall, for example, UNESCO's depiction of professional journalists as “communicators who work in the service of truth,” a ringing endorsement of the democratic and epistemic importance of practitioners of that profession. Elsewhere, in a guide to combating what they call democracy's current “information disorder,” Christine Wardle and Hossein Derakhshan write that “initiatives to help build trust and credibility” in the legacy news media “go hand in hand with any initiatives aiming to combat mis- and dis-information.”¹⁰⁷ Here, again, the legacy news media and professional journalists are seen as key players in the drive to secure democracy's epistemic well-being. These attitudes are exceedingly common for Progressives.¹⁰⁸ From this it may be inferred that at least one *kind* of information Progressives believe citizens need to be able to access and demonstrate a propensity to place their trust in is the kind of information disseminated by professional journalists, i.e. the legacy news.

The second commonality shared by most commentators who view 21st century technology as a threat to the epistemic well-being of democracy is a belief that democracy is best off when people place their trust in the scattered array of researchers, technicians and practitioners known collectively as the “scientific community,” rather than those who contest their claims. Lewandowsky and colleagues, for instance, paint as “dystopian” a future in which citizens have “had enough of experts” and lament especially that among certain citizens there has been a “general decline of trust in science.”¹⁰⁹ A similar view is espoused by Anya Schiffrin, who names “institutions of education, science, and media” as forces that “have traditionally served to keep 'false facts' and demagoguery at bay.”¹¹⁰ Again, such views of the “scientific community” and the importance of citizens placing their trust in the claims it disseminates are par for the course among Progressives.

¹⁰⁷Wardle and Hossein (2017), 63

¹⁰⁸See, for instance, Deb, Donahue and Glaisyer (2017).

¹⁰⁹Lewandowsky et al (2017), 354-8

¹¹⁰Schiffrin (2017), 123

This respect for the legacy news media and the “scientific community” and a correspondingly high estimation of the reliability of the information disseminated by them gives us a way to connect the Progressive Ideal to real-world phenomena in a way that might be used to gauge the epistemic well-being of democracy. That is, it gives us *standards* we can use to assess how well democracy is doing epistemically. These standards are as follows:

Progressive Standards for Assessing the Epistemic Well-Being of Democracy

1. Citizens ought to be made aware of the kinds of information generally disseminated in the legacy news.
2. Citizens ought to be made aware of, and willing to place their trust in, the claims of professional journalists and scientists, as well as of the institutions of professional journalism and of science.

Of course, this is nowhere near a comprehensive list of the standards the informational environment would need to meet in order for democracy to be doing well epistemically according to the Progressive Ideal. I have no intention of developing such a list. All I mean to do here is to say that the two standards provided above are *among* those that are implicit in the views commonly aired by proponents of the Progressive Ideal.

Criteria for Research About the Epistemic Well-Being of Democracy

The two standards listed above give us a connection between the abstract principles articulated Progressive Ideal and the real world. They offer us a means of assessing how well democracy is doing according to the Ideal, at least when it comes to the relationship between the public, scientists and professional journalists. Until we actually *perform* such an assessment of our informational environment's facilitation or obstruction of citizens' efforts to meet the standards the Progressive Ideal requires, though, no reliable verdict can be issued about the epistemic well-being of democracy. To make such an assessment requires research. What might that research look like?

I believe there are four criteria that must be fulfilled if research is to give us the ability to confidently issue a verdict about whether 21st century technology, or any other phenomenon, is

negatively impacting the epistemic well-being of democracy. First, and most obvious, the research must be *rigorously and appropriately empirically informed*. Once explicit standards have been articulated, like the two I articulated in the previous sub-section, the question becomes not *which* standards ought to be embraced but *whether* our current informational environment meets those standards. This is an empirical question. Any time we are faced with an empirical question, the first order of business is to conceive of and carry out the best empirical research we can think of to give us insight on the matter. While acknowledging limitations both cognitive and material, we ought to try and make sure our analyses are based on the *most rigorous* empirical data available and that research designed to assess and cited as evidence of the epistemic condition of democracy is *appropriately designed* to capture all the important aspects of that epistemic well-being.

The second criterion that should be fulfilled if research is to lead to a clear verdict about the epistemic well-being of democracy is that such research must *articulate the counterfactual standard* against which democracy's current epistemic well-being is being compared. No serious thinker believes humans have unlimited powers of cognition or infinitely dependable stores of benevolence. What is attainable by actual humans living in the real world is often the stuff of despair when compared against the things we can imagine in our minds. When Augustine compared the world into which he had been born with the vision of paradise his mind was able to produce, so great was his disappointment at the state of terrestrial affairs that he lost hope for it, placing them instead in the hope of a future life to live up to the vision in his mind. While I find Augustine often to be a captivating writer, I must admit that I have little admiration for this sort of response to nature's long-standing refusal to mold itself to our wishes and imaginings. Political theorists do their job best not by producing visions of paradise and condemning reality when it falls short of their fancies, but by figuring out ways for us to make the most of the crooked timber¹¹¹ with which we have to work. When it comes to issuing verdicts about the epistemic well-being of democracy, then, our standard of comparison ought not to be some Augustinian fancy but instead a state of affairs for which a defensible case can be made that humans as we know them (or as we might reasonably expect them to become) are capable of

¹¹¹This phrase is, of course, inspired by Kant's "Idea for a Universal History with a Cosmopolitan Purpose" (Kant 1991 [1784])

attaining it. In the current case of whether 21st century technologies are detrimental to the epistemic well-being of democracy, that means that we cannot simply identify a litany of what we take to be sub-optimalities in the informational environment and jump to the conclusion that technology is harming democracy epistemically. Instead, we need to *compare* the current state of epistemic affairs against the world that we think *would* exist without those technologies, and see how the former fares compared to the latter.

The third criterion that must be met by research seeking to assess the epistemic well-being of democracy is that it must take into account *both* of our epistemic aims, which are to (1) *discover truths* and (2) *avoid mistaken beliefs*. These two goals correspond to the “Type 1” and “Type 2” errors familiar to empirical researchers. Considerably more contemporary social science literature has been dedicated to warning of the dangers of *misinformation* online, a case of Type 2 error, than has been written about the danger of failing to enable citizens to maximize their potential for adopting true beliefs, which, since it involves a failure to embrace truth, is a corollary of Type 1 error. While avoiding the proliferation of misinformation is clearly important, avoiding mistaken beliefs is equally clearly not all we should want from citizens. This has perhaps been most effectively illustrated in recent decades by feminists, gender theorists, critical race theorists, and others who have brought to light the perspectival limitations of privilege and the need to therefore proactively emphasize the inclusion of diverse perspectives into decision-making bodies and procedures. As these thinkers have shown us, when it comes to collective knowledge, we want democracy to be able to access and respond to *lots of truths* from *lots of perspectives* just as much as we want it to *avoid* being responsive to (what are currently taken to be) *untruths*. Research that is cited to subsidize conclusions about the epistemic well-being of democracy needs to take both of our epistemic goals – to *discover truth* and to *avoid mistaken beliefs* – into account.

The final criterion whose importance I want to emphasize for work examining the epistemic well-being of democracy is that its claims needs to *account for causal complexity*. The internet and social media are massive entities and their effects on society manifold. The number of web *sites* (just domains, not total pages) on the internet was recently estimated at over a

billion.¹¹² The number of social media users is thought to be upwards of three billion.¹¹³ These sites vary enormously in subject matter, user base, degree of sophistication, and a host of other variables, as do the ways in which particular users interact with them. It is very likely that the effects of social media, web forums, the “blogosphere,” and the internet as a whole vary enormously both *between* and *within* individuals. Given the complexity of this situation, it is doubtful that generalizing about the *overall* impact of “the internet” and “social media” on democracy is the most helpful way to discuss 21st century technological changes' effect(s) on democracy. One reason this is so, which I've already hinted at and to which will dedicate an entire chapter of this dissertation, is because individual-level behaviors often produce counter-intuitive aggregate outcomes.¹¹⁴ Another, which I will address in this chapter, is the ways in which sources deemed by pessimists about the epistemic impact of the internet and social media as reliable, such as professional journalists and scientists, *themselves* depend upon 21st century communications technologies both when it comes to forming their own knowledge and when it comes to disseminating their claims to the public. Researchers making claims about the negative impact of 21st century communications technologies on the epistemic well-being of democracy and the solutions we ought to consider to address them need to take this causal complexity into account, otherwise they risk encouraging attitudes in citizens and supporting public policies whose effects are worse than the problems they are trying to combat.

These, then, are the four criteria research needs to fulfill if it is to pronounce reliable verdicts about the epistemic well-being of democracy:

¹¹²Huss (2021)

¹¹³Tankovska (2021)

¹¹⁴This principle was perhaps most famously summed up by Adam Smith, who pointed out that it “is not from the benevolence of the butcher, the brewer, or the baker, that we expect our dinner, but from their regard to their own interest” (Smith 1937 [1775], 14). That passage highlights that what may seem like, and in fact *be*, completely self-interested behavior on the part of the individual agent may nevertheless produce socially beneficial effects when generalized across the population. Similar arguments have been made in philosophy of science, both with regard to the role of the “priority rule” of bestowing the financial and reputational awards of discovery upon the first discoverer (Strevens 2003) and with regard to scientists' reluctance to expose themselves to (Weisberg and Muldoon 2009) or update their belief in response to (Zollman 2010) new findings in their field that contradict their own beliefs.

Criteria for Research on the Epistemic Well-Being of Democracy

1. Must be based in rigorous and appropriate empirical evidence
2. Must articulate the counterfactual against which the current epistemic state of democracy is being compared
3. Must take into account both epistemic commandments – *discover new truths* and *avoid error*
4. Must account for causal complexity

While these criteria may not be *sufficient* to diagnose the epistemic well-being of democracy according to the Progressive Ideal, they are *necessary* to that end. Research that does not *at least* fulfill these criteria simply cannot serve as the grounds for issuing a verdict about the epistemic well-being of democracy.

In the rest of this chapter, I will argue that none of the research and commentary that characterizes 21st century technologies as threatening to the epistemic well-being of democracy fulfills all these criteria. On the contrary, I will argue that, given the current state of the evidence, due consideration of all four criteria may as easily lead to the conclusion that the epistemic well-being of democracy is not only not declining, but is in fact thriving, and it is doing so at least in part *because of* the same 21st century communications technologies that are so often blamed for its decline.

Section Three: An Immanent Critique of Progressives' Interpretation of the Impact of 21st Century Communications' Technologies' Epistemic Impact on Democracy

In this section I will show examples of scholarly and lay neglect of each of the criteria given above. The result of this neglect, I will argue, is not only to cast doubt on the verdict that the internet and social media are harming democracy's epistemic well-being (at least compared to any plausible counterfactual scenario), but to open the door to the possibility that exactly the opposite is the case – that the internet and social media are in fact helping the epistemic well-being of democracy to flourish (again, at least compared to any plausible counterfactual), or at least mitigating the harm that it would be suffering otherwise. This in turn calls into question

Progressives' deeply pessimistic view of the judgmental capacities of lay citizens as well as their deferential model of civic epistemology.

Progressive Neglect of Criterion 1: Is a cluttered informational environment disproportionately unfavorable to science and professional journalism?

One claim frequently forwarded by Progressives is that 21st century communications technologies are comparatively unfavorable to scientists, professional journalists and similar experts while being comparatively favorable to unreliable sources of information. Anya Schiffrin, for instance, characterizes the internet as “a forum for distributing information that does not adhere to typical¹¹⁵ standards of truth, scientific inquiry, and evidence-based news and information.”¹¹⁶ It is clear that by “typical” Schiffrin means typical *of* the types of institutions she views with favor, i.e. the legacy news media and the “scientific community,” and it is clear from the rest of her essay that she views the legacy news media and the “scientific community” as typically producing “evidence-based news and information.” Elsewhere, the British House of Commons Report on Disinformation and “Fake News” warns that the consequences of people “increasingly finding out about what is happening...through social media, rather than through more traditional forms of communication” like television and newspapers could be “devastating” to democracy.¹¹⁷ Clearly, the committee drafting the Report views “more traditional forms of communication” as friendlier to reliable information than social media, and I do not think it is much of a stretch to suppose that a significant part of the attraction of “traditional” forms of communication was that they were dominated by professional media organizations, many of which, like the *BBC*, had highly respected and highly professionalized news branches. In portraying the internet and/or social media as unfavorable environments for the dissemination and acceptance of expert and professional sources of information and favorable environments for

115 Note that although Schiffrin uses the word “typical,” she never explains which community, group or era that term is meant to refer to. Its use seems to imply that Schiffrin believes the internet to be inimical to both scientific inquiry and evidence-based news and information, compared against some standard that we should *normally* expect to prevail. However, she never describes what that normal state of affairs is or why we ought to share her expectation that it would normally apply. She thus fails to fulfill Criteria 3 for research seeking to render a verdict on the epistemic well-being of democracy, which is to articulate the counterfactual standard against which the current state of affairs ought to be compared.

116Schiffrin (2017), 123, my emphasis

117House Report (2018), 4

less reliable sources, Schifffrin and the House Report writers exhibit a tendency that is very widespread in the current moment, both among Progressives and among the populace in general.¹¹⁸

Widespread though it may be, however, this supposition is backed by very little empirical support. Rather than basing their claims in rigorous and appropriate empirical research, commentators typically either cite evidence that is anecdotal or inconclusive, make assertions unbacked by empirical evidence, or – most commonly – simply *assert that* the internet has (apparently self-evidently) fueled a widespread abdication of scientific and journalistic expertise. But for each of the most commonly forwarded Progressive lines of argument in this vein, I believe that further consideration of the available empirical record reveals there to be as much support for optimism as there is pessimism. For instance, Progressives often make much of the fact that social media's barriers to publication are significantly less formidable than those met with by public affairs prognosticators in the pre-digital era. Where once the written or spoken word had to be curated by book and journal publishers, television or radio show producers, or a newspaper's editorial board, the internet and social media allow practically anyone to publish thoughts that can be seen by the public in real-time, without the need to go through any such curatorial process. This development is supposed by Progressives to have led to an informational environment where the voices of officials, accredited experts, professional journalists and the like carry diminished weight with the public, as they now have to compete with an influx of competition from amateurs, political opportunists, automated bots, and other (presumably) less-reliable sources. The idea seems to be that the public is unequipped, either epistemically or logistically, to navigate this cacophony of voices and find their way to the reliable ones. Because of this, Progressives fear that large numbers of democratic citizens in the internet age are abandoning authorities who ought to be trusted in favor of those who oughtn't, giving up the effort to find the truth at all in a fit of exhaustion, or some combination of both.¹¹⁹

The logic behind these worries is intuitively appealing. After all, the knowledge gap between experts and lay persons makes it difficult for the latter to figure out which of the former to trust in the first place; indeed, the problem is so vexing that epistemologists disagree about

¹¹⁸See, e.g., Jackson (2017), Nichols (2017), Deb, Donahue and Glaisyer (2017)

¹¹⁹Versions of this line of reasoning can be found in Davies (2016); Deb, Donahue and Glaisyer (2017); and Lewandowsky et al (2017)

whether it can be epistemically responsible for them to do so at all.¹²⁰ Increase the number of voices with which the experts have to compete and that task seems bound to become more difficult.¹²¹ Add this to this the long-standing and, in my view, compelling argument forwarded by political scientists for over half a century that the meager rewards of voting are insufficient to motivate most citizens to exert much effort in order to become informed¹²² and you get a somber view of the epistemic prospects of a society served by 21st century communications technologies indeed. By making the already daunting and unmotivating task of acquiring sound political knowledge even more burdensome by increasing the volume of competing voices the average citizen has to wade through to get to it, the flood of voices let loose by the internet and social media might persuade even *more* of them that the expenditure of time and effort simply isn't worth it, thus diminishing the already worryingly low levels of political knowledge among democratic publics.¹²³

Intuitive though this logic may be, however, the proposition that the proliferation of sources made available by the internet and social media's extraordinarily low barriers to publication has resulted in either a diminishing of the epistemic influence of reliable vis-a-vis unreliable sources of information within society or an increased propensity for citizens to abandon the quest for truth out of epistemic exhaustion or frustration are empirical claims. They

120Few have argued that it is *always* wrong for lay citizens to defer to experts but, as Linda Zagzebski argues, the absence of such a position among epistemologists is strikingly at odds with their near-universal endorsement of epistemic egalitarianism and epistemic self-reliance (Zagzebski 2012, ch. 1). Among those who do (perhaps self-contradictorily) embrace these normative principles there is considerable disagreement about how willing lay persons ought to be to defer to the testimony of experts. At least since John Hardwig (1985) pointed out in a seminal article how, contrary to the presumptions of many epistemologists up to that point, even scientists are heavily dependent on expert testimony, most epistemologists seem to believe that such deference is very often justified in contemporary life. The differences mainly revolve around how readily citizens ought to be willing to do so and under what conditions that deference is justified. For a synopsis of the various views on the subject, see Lackey (2011).

121Minozzi (2010) seems to defend a similar argument in his “jamming” theory of politics, though his focus is on *intentional* efforts by malicious actors to deter citizens from identifying and placing their belief in the expert consensus. The argument typically forwarded by those who are worried that the decrease of publication costs is likely to produce an informational environment that is especially unfavorable to scientists, professional journalists and other experts vis-a-vis less reliable sources would extend Minozzi's theory beyond such intentional efforts at obfuscation. They seem to be saying that the *mere* increase in ease of publication of and access to information by an increasing variety of sources is likely to undermine the societal influence of scientists and professional journalists (etc.).

122This is the logic behind Bryan Caplan's characterization of voters as “rationally irrational” (Caplan 2007).

123That democratic citizens *are* largely uninformed is, of course, one of the most widely supported conclusions empirical findings of public opinion researchers (e.g. Berelson et al 1954; Converse 1964; Delli Carpini and Keeter 1996; Somin 2016).

are best assessed not by the intuitiveness of the causal stories that spring to our minds on considering them but by whether or not they are supported by the best available evidence. *Does* the empirical record suggest democratic citizens are turning from reliable to unreliable sources or giving up the quest for truth entirely? Let us deal with each of these propositions in turn, beginning with the claim that lowered costs of publication have prompted people to turn away from reliable sources of information and to flock to less reliable information sources. There are numerous assumptions built into the various arguments that have been raised in defense of this claim. One commonly made assumption would appear to be that *as a general rule* curation is likely to *increase* the reliability of the claims that are curated, compared with information that is not. At first blush, this assumption makes some sense. After all, the whole point of *having* a process of curation is often thought to be to *improve* the quality of the information that makes it through that process.¹²⁴ Reasonable as this may seem in the abstract, however, there is no necessary connection between curation and an increase in the reliability of a claim. Take, for example, the curatorial process employed by the legacy news media. Throughout the history of publishing right up to the present day, at least *some* publishers at *all* levels of professionalism have taken quotes out of context, spliced video clips in a way that produces a misleading message, reported only selective portions of a story, and reported only stories whose content would seem to support a given worldview. This may happen for any variety of reasons – the personal agenda of reporters, editors or executives at a given legacy news organization; the monetary incentive to appeal to (or at least not offend) the ideological makeup of a news organization's audience or its advertisers; the logistical need to fit stories into compact spaces and time windows; etc. Whatever the reason, the upshot is that while the curatorial process used by the legacy news media has often worked to “reduce the number of errors that might otherwise be reported,”¹²⁵ in the words of epistemologist David Coady, sometimes it has worked to no beneficial epistemic effect at all, and sometimes to a negative one. As Coady points out, some reporters' careers have “flourished despite (and sometimes, it appears, because of) having repeatedly reported falsehoods (or repeatedly being reckless with respect to the truth of what

¹²⁴See Goldman (2008) for a prominent epistemologist's articulation of this view.

¹²⁵D. Coady (2011), 289

they report). What is more, there are reporters who appear to have been sacked (or had their careers stall) for reporting the truth.”¹²⁶

While the legacy news media and the “community of science” are not shy about publishing the *ideals* that, in theory, guide their publication process, as I will explain in a future chapter, in the vast majority of cases the public has no ability to tell with what degree of rigor the curatorial processes *claimed* by these and other institutions like them are *actually* followed. Recognizing this difficulty, some will prefer to base their estimations of the accuracy of legacy news reports on analyses of incentive structures to anticipate the *likely* behavior of news institutions vis-a-vis the types of sources – bloggers, social media users, “citizen journalists,” etc – that have emerged in the internet era. In fact, many already do take this angle, citing structural incentives like a news company's interest in maintaining a reputation for reliability, the existence of competitors vying for the public's trust, and so forth as reasons why we ought to trust such institutions more than the types of sources favored by the internet and social media. For instance, Hunt Allcott and Matthew Gentzkow write that “reputational concerns discourage mass media outlets from knowingly reporting false stories.”¹²⁷ While such incentives may well exert pressure on mainstream news sources to stick closer to the truth, it is not clear that this single incentive typically *outweighs* other incentives, such as the need to grab viewers' attention, confirm the anticipated audience's worldview, and satisfy the personal and/or political agendas of those in administrative positions within that news organization. Nor is it clear that similar incentives do not apply to non-legacy news sources on the internet. Why should the non-legacy news sources that have emerged in the internet era care any less about retaining viewers over the long-term than legacy sources? Both depend on gaining and sustaining the attention of a regular viewership for advertising revenue and therefore profits and presumably if people are likely to punish one type of information source for unreliability they will be just as ready to do so for another. If both are subject to similar incentives, why we should assume that *only* legacy news sources experience are motivated by those incentives to publish truths and maintain a reputation for trustworthiness while 21st century rivals for the public's attention are not and do not? It is not clear that the logic that lies behind Progressives' implication that the incentives faced by legacy

¹²⁶Ibid., 281

¹²⁷Allcott and Gentzkow (2017), 214

news companies are likely to keep them honest while the same does not apply to internet-era newcomers is sound.

Nor is it obvious that it is supported by the empirical record. One of the articles that is most widely cited by Progressives as evidence that unreliable sources are coming to rival reliable ones in the internet age is Craig Silverman's *Buzzfeed* article that purports to demonstrate how the top 20 fake news stories outperformed the top 20 legacy news stories in the months immediately preceding the 2016 Presidential election. Even if we take at face value Silverman's categorizations of “real” vs. “fake” news (which I largely agree with) and ignore the substantial methodological problems with the article,¹²⁸ a comparison of the fate of the web domains reporting what Silverman categorizes as “fake news” versus the fate of those reporting the “real news” suggests that citizens are likely to punish purveyors of faulty information and reward those who give them the “good stuff.” Of the 13 fake news websites Silverman cites as having supplied the most popular fake news stories between September and November of 2016, 10 of them are now dead URLs, one (*The Burrard Street Journal*) is an established satirical publication along the lines of the *Onion* dedicated mostly to sporting “news,” one is Breitbart News, most of whose contents are more accurately characterized as heavily conservatively biased than “fake news,” and the other is the *World News Daily Report*, whose estimated monthly traffic is between 0.5-1.5 million viewers.¹²⁹ To put things in perspective, that is less than 0.5% of the estimated daily traffic of *The New York Times* alone.¹³⁰ Meanwhile, not single one of the 20 sites listed by Silverman as “real news” sources has gone defunct. On the contrary, many of them continue to rank in among the top-500 most trafficked websites in the world. That so many of Silverman's fake news sites have collapsed in the short space of three years while the web traffic of “real news” sources continues to boom suggests lay citizens *are* rewarding what Silverman and other Progressives view to be reliable sources of information and punishing those that do not.

¹²⁸Briefly, the two major problems are (1) Silverman's lumping-in of both editorials and articles purporting to be factual reports and, more importantly, (2) his use of a *single source's version* of a story as an estimate of that story's “influence” on social media, rather than including *all* news source's reports of the same incident(s). The latter of these is likely to severely underestimate the extent to which accurate stories of significant political importance are spread on social media vis-a-vis fake news stories.

¹²⁹<https://www.similarweb.com/website/worldnewsdailyreport.com#overview>

¹³⁰<https://www.similarweb.com/website/nytimes.com#overview>

What about the curatorial processes used by our other working example, the “community of science?” Are these processes any more guaranteed to improve the reliability of the knowledge published by that “community’s” members than those employed by the legacy news media? While I personally believe the curatorial processes used by the “community of science” to have historically been largely epistemically beneficial, I also believe they *can be*, and sometimes *have in fact been*, used to hamper, rather than aid, the pursuit of truth. For instance, confirmation of the causal role played by *H. pylori* bacteria in causing ulcers was set back by perhaps more than half a century by the persistent refusal of the gastroenterology community to consider any hypothesis attributing ulcers to bacteria. Indeed, so adamant was that community’s refusal to consider the hypothesis that a Greek doctor, John Lykoudis, was refused publication on and eventually fined for treating ulcer patients with antibiotics in the 1950s and 1960s, despite the fact that “by all reports he was very successful.”¹³¹ The *H. pylori* hypothesis would not be entertained by the gastroenterology community until the early 2000s, when researcher Barry Marshall became so desperate that he drank a vial of the bacteria himself, after which he immediately “became ill and was able to cure himself with antibiotics.”¹³² I personally doubt that the curatorial processes employed by professional journalists and scientists are usually as dramatically opposed to the discovery of truth as gastroenterologists’ treatment of peptic ulcer research was in this instance. However, such instances demonstrate that *merely* going through a curatorial process carries no necessary implication for the quality of the information that makes it through that process. Nor should this conclusion surprise us. After all, tabloids like *The National Enquirer* and conspiracy publications like *Infowars* also extensively curate their content, yet few would say that this has always, or even usually, improved the reliability of the claims that passed through that process. If anything, many would argue the opposite is the case – that the curation process for these and similar entities causes them to weed out the comparatively boring truth in favor of spectacular falsehoods. We cannot, then, leap from the observation that a decreasing proportion of the information published within a given public at a given time has been subjected to a process of curation to the conclusion that that information must therefore, on average, be less

¹³¹Zollman (2010), 20

¹³²*Ibid.*, 21

reliable. To do that, we need to know more about the *nature* of the curatorial processes typically being employed as well as the *actual use* to which they are being put.

These considerations cast doubt on the frequently made assumption that the curatorial processes claimed to be integral to the functioning of knowledge-disseminating institutions like professional journalism and science are bound to make the knowledge they disseminate more reliable than alternatives. But even if they didn't, we would still not be justified in concluding that upticks in the proportion of non-curated to curated information in democratic citizens' informational environment are bound to harm the epistemic well-being of democracy. This is because, contrary to the expectations of pessimists about the epistemic impact of 21st century communications technologies, upsurges in the availability of sources of information castigated as unreliable by the learned of their era have historically *not* led to a decline either in the sophistication or societal influence of science or in the development and increasing societal influence of higher-quality journalism. If anything, the opposite has been the case.

The number of publishers and the volume of published matter have grown exponentially throughout the history of contemporary representative democracy. Throughout this time, learned commentators have regularly voiced indignation about the generally deplorable quality of the publications that seemed suddenly to be flooding into their informational environment. During the early days of the American republic, Thomas Jefferson wrote, "Nothing can now be believed which is seen in a newspaper. Truth itself becomes suspicious by being put into that polluted vehicle."¹³³ When, a few decades later, new printing technologies enabled the mass publication of the penny papers, they were castigated by the respectable of their era as full of "indecent, blasphemy, blackmail, lying, and libel."¹³⁴ Steady improvements in printing technology allowed even more information to flood the 19th century informational environment so that between 1790 and 1915, the number of newspapers per inhabitant of the United States increased from about one per 25,000 in New England and the North Atlantic to about one per 7,000.¹³⁵ Toward the end of the century, scholars were no more optimistic about the public's ability to navigate this changing informational landscape than was Jefferson at the end of the previous. Instead, they called their age the age of "yellow journalism," viewing its most popular publishers as unreliable

¹³³Jefferson to Norvell (1807)

¹³⁴Schudson (1978), 55

¹³⁵Dill (1928), 77

and unprincipled sources of information that were imperiling democracy. In the first half of the 20th century, thinkers like Walter Lippmann worried about what would happen to a democracy served by such papers as well as the newly emergent forms of propaganda and public relations spread both by paper and by radio. For a brief period thereafter, the number and variety of Americans' favored sources of information was constricted by the dominance of broadcasting technologies, which required capital-intensive industries in order to disseminate messages widely. However, when cable television became a staple of American households in the 1980s and 1990s, communications and political science scholars filled many a journal with worries about declining newspaper subscriptions and nightly news ratings as viewers turned instead to entertainment alternatives typically considered to be of inferior epistemic merit.

The point to note here is that with the possible exception of the broadcast era, the history of contemporary democracy is a history in which the *number* and *types* of sources available to citizens has been *increasing*. If the logic endorsed by many who have warned of the threat posed to epistemic well-being of democracy by the decreased costs of publications made available by the internet and social media were an accurate depiction of the general relationship between the quantity of (what scholarly contemporaries judged to be) low-quality information available to lay citizens and their propensity to embrace high-quality sources, given the persistent derision and suspicion with which well-respected thinkers in every era have viewed the information disseminated by new entrants into the publications market, one would expect that each era should have seen a *decrease* in the overall societal influence of science and high-quality journalism. After all, each new influx of publishers and each new uptick in the amount of low-quality information with which the average citizen was beset should have made it *more difficult* for lay citizens to avoid falling prey to the claims of the silver-tongued charlatans, propagandists, and other purveyors of low-quality information that, supposedly, are best poised to take advantage of lowered barriers to publication. Yet, this expectation appears to be entirely confounded by the historical record. The time period stretching from Jefferson's fulminations against the press to the present is one in which the sophistication and societal influence of science have increased consistently and rapidly. Indeed, no group during that period experienced so meteoric a rise in the public's estimation of its epistemic reliability. During the period between Jefferson and now, practitioners of science have gone from a band of tinkering, eccentric

aristocrats who published in obscure journals and occasionally managed to impress the domestic political authorities to a highly professionalized class of knowledge-seekers who are viewed by the majority of the public as the bearers of the most reliable knowledge in a wide range of fields. Until very recently, a similar story was true for the sophistication and societal influence of high-quality journalism. The eras whose regular influxes of new entrants into the publications market drew such consistent ire from intellectual contemporaries saw the most popular newspapers develop from ad-hoc collections of whatever reports, advertisements, fictions or ill-tempered insults of public figures found their way to the local printer that day (Colonial era) to publications tied explicitly to political parties (early 1800s) to institutions employing several on-the-ground reporters with regular beats and an ethos of neutrality (early-late 1800s) to large-scale institutions employing specialized reporters (late 1800s-early 1900s) to professionalized institutions with formalized standards, procedures and training (mid 1900s-present).¹³⁶ All through this time, paper newspaper subscriptions continually increased until the 1990s, when they began their now-infamous decline.¹³⁷

In the case of both the “community of science” and journalism, in other words, both the sophistication of the methods employed by practitioners and the public's estimation of by them have *consistently increased* over the last two-and-a-half centuries, even though all through that era the number of publicly available sources of information whose epistemic reliability was disparaged by the learned of the day *also* consistently increased. While I am not a believer in the inevitable triumph of truth and reason and do not think the past two-and-a-half centuries are evidence of that triumph's inevitability, I do believe that the historical record demonstrates that whatever the relationship (if any) between decreased costs of publication, the proliferation of information sources considered to be of poor quality by contemporaries of the day, and the sophistication and societal influence of reliable sources of information, it is considerably less straightforward than is commonly assumed by Progressives. As such, if we are to estimate the epistemic impact of 21st century technologies on democracy's epistemic well-being, we must do

¹³⁶Schudson (1978)

¹³⁷I say “decline” even though, due in part to differing views on what ought to count as “high-quality” journalism and due in part to the difficulty of tracking digital subscriptions (as well as news readership by those who do not carry a paying subscription to a major periodical), it is not entirely clear that the reading audience for news journals has declined at all, either in terms of actual eyes-on-screen readership or in terms of the effective influence of the written news industry.

more than rely on such intuitive chains of reasoning. Instead, we must ground our analyses in careful and rigorous empirical research. Up to now, Progressives have not done so, and the contemporary and historical examples given here suggest that if they did, the results may surprise them.

What of Progressives' other worry, that citizens inhabiting the 21st century's informational environment might be so daunted by the scope of available information that they will abandon the quest for politically relevant truths entirely? Again, the reasoning behind this supposition might at first seem intuitively plausible, since the proliferation of information in the internet era can seem overwhelming and it seems reasonable that people who are overwhelmed by alternative sources will be more likely to give up the search for truth than those whose search is less difficult. But the closer one looks at this line of reasoning, the more suspect it becomes. If the descent into epistemic nihilism was a common response to the proliferation of sources of information, one would surely expect entry into any large library, with its countless shelves of volumes on any number of topics, to have sent large quantities of the public into epistemic nihilism long since. Yet, few Progressives argue against the proliferation of books, journals and magazines of the sort that typically fill such institutions; on the contrary, libraries are typically viewed as important *contributors* to public knowledge. The very idea that high volumes of information typically have an inverse relationship with potential inquirers' interest levels in that knowledge should be suspect in the eyes of any scholar. Nothing is more common for the scholar than to find that the topic they are studying is vastly more complicated than she first considered, that more has been written on it than she ever dreamed, and that the truth is far more difficult to decipher than she initially thought. Yet, these things can as easily serve to *motivate* and *captivate* researchers as cause them to despair. Surely these responses to unanticipated complexity are just as available to the lay public as they are to researchers, so it is not clear that *just because* the internet provides more sources of information than ever we should expect citizens to therefore become less motivated to find the truth.

Moreover, the historical record once again seems to contradict the line of argument forwarded by Progressives. Each new upsurge in the amount of information and breadth of sources made available to the public by improved technologies and each move toward the liberalization of policies regarding freedom of speech and the press for the past two and a half

centuries has been accompanied by increased literacy rates, improved education levels overall, an increase in the number of people dedicating a greater number of the years of their life to the pursuit of education,¹³⁸ and an increased number of total consumers of the voluminous informational resources emerging into the informational environment.¹³⁹ The consistency of this relationship between increases in the information made available to the public and increases in the public's efforts to make use of that information by becoming educated in various ways has been well-documented by historians.¹⁴⁰ Far from prompting citizens to bury their heads in the sand in an effort to drown out the overwhelming cacophony of voices unleashed to them by improved communications technologies, the historical trend has been for the public to *increase* their pursuit of literacy and education in eras during which information has become more widely available. Of course, this doesn't necessarily mean citizens were driven to new heights of curiosity *by* the increased availability of information, or even that it was curiosity, rather than material necessity or public policy, that drove them to become literate or pursue an education. However crude they may be as metrics of the lay public's motivation to seek the truth they may be, though, looking at literacy and education rates *at least* makes an attempt to ground our analysis in some kind of empirical evidence rather than an intuitive-seeming logical chain or a handy anecdote, both of which are too often resorted to by Progressives arguing that 21st century communications technologies pose a threat to democracy. And its implications are just the opposite from that supposed by Progressives.

I believe that a consideration of these factors casts serious doubt on the supposition that a proliferation of information sources is likely to lead large numbers of citizens to abandon the quest for truth entirely or, even more implausibly, “stop believing in facts altogether.”¹⁴¹ Again, it

¹³⁸Roser and Ospina (2019)

¹³⁹William A. Dill estimated in 1918 that the total circulation of newspapers in all states increased from around 4.5 million in 1850 to around 19 million in 1872. By 1940, Pew reports total weekday circulation to have eclipsed 40 million (Barthel 2017), joined by that time of course by growing radio and television news audiences. The number of people engaging in a good-faith effort to seek out truths on the internet is, of course, very difficult to assess, but consider that the number of subscribers to the TED Talks Youtube channel alone is currently over 14 million, the number of people who log on to nytimes.com each day numbers around 10 million, and the number of people who “like” physicist Neil deGrasse Tyson's Facebook page – and will correspondingly be alerted when he posts – numbers over 4 million.

¹⁴⁰Starr (2004)

¹⁴¹The quote is from Lewandowsky and colleagues, who cite two studies in support of this remarkable claim. But both studies they cite merely demonstrate that subjects presented with both scientific evidence and misinformation in an experimental setting fail to immediately embrace the accepted scientific position (Lewandowsky et al 2017, 355). This is very far from demonstrating that an excess of available information

is *possible* that there is a general tendency for upsurges in available information to prompt people to give up searching for the truth despite these historical trends in literacy and education. It is also *possible* that our era is different than these prior eras, such that the quantity of information made available to citizens by the internet and the speed at which it is generated and transmitted are finally great enough that people feel overwhelmed rather than inspired by it. However, these are matters that cannot be settled by intuition and anecdote; they require rigorous and appropriate empirical study. In the meantime, we ought to recognize that, given the tendency for similar concerns voiced by intellectuals in prior eras to have proven ill-founded, there is no reason for commentators to err, as they currently seem to do, on the side of pessimism. If anything, history gives us reason for hope on this front. The overall point is that the relationship between proliferating information sources and (a) citizens' tendency to favor reliable vs. unreliable sources as well as (b) citizens' motivation to seek the truth are complicated empirical questions that need to be addressed via well-conceived and careful research, and that the available empirical and historical evidence is very far from giving clear support to the Progressive position. Until a more rigorous empirical case is made to explain why, in apparent contradiction with historical tendencies, we should anticipate a proliferation of information sources to result in either declining accuracy in citizens' beliefs or a large-scale public abandonment of the quest for truth, we have no good reason to follow Progressives in reaching such pessimistic conclusions about the prospects for lay citizens to learn to effectively navigate the 21st century's changing informational environment or for the epistemic well-being of democracy generally.

Before turning to the next section, I want to make one more argument. I believe that one major reason Progressives' interpretation of the epistemic impact of 21st century communications technologies on democracy is so consistently pessimistic despite the inconclusiveness of present evidence and that pessimism's lack of historical grounds is that Progressives *presuppose* a pessimistic view of lay citizens. Specifically, they presuppose that lay citizens' *capacities of judgment* are insufficient to help them recognize truth and avoid error by the lights of that judgment. This presupposition makes sense of Progressives' pessimism in an elegant and straightforward way. If we have a *deeply pessimistic* view of the judgmental capacities of lay

keeps people from seeking out the truth, much less embracing the sort of radical skepticism implied by the authors and others who have reacted the most volubly about the onset of the so-called "post-truth condition" (e.g. Calcutt 2016, Williams 2017).

citizens, the logic of Progressives makes sense. If, for example, we believe citizens' judgment is poor enough that they are *as* or *more* likely to believe falsehoods than truths and to place their trust in unreliable rather than reliable sources of information, then it makes perfect sense that we would expect the epistemic well-being of democracy to decline as the quantity of informational sources easily accessible by the public increases. After all, each new source of information is an extra opportunity for the lay public to make a judgment, and if their judgment is very poor their mistakes are likely to outnumber their successes. As I showed in the introduction and opening chapter to this dissertation, this deeply pessimistic view is entirely consistent with many Progressives' depictions of the lay public, and I believe the adoption of this deeply pessimistic view is a major reason causes Progressives so frequently suppose that increasing the number and variety of sources of information easily available to the public is likely to lead to bad collective epistemic outcomes. What I want to point out here is the way this deeply pessimistic *starting point* tends to *reinforce* Progressives' pessimism on other fronts. It is precisely *because* Progressives *begin* with a deeply pessimistic view of the judgmental capacities of the lay public that Progressives *suppose* 21st century communications technologies will produce, and are currently producing, negative collective epistemic outcomes. This supposition casts a negative hue on Progressives' interpretation of the democratic impact of 21st century communications technologies, which in turn reinforces their pessimism about the judgment of the public.

I hope to have shown that it is far from clear that either form of pessimism – Progressives' pessimism about the judgmental capacities of lay citizens or their pessimism about the epistemic impact of 21st century communications technologies on democracy – is warranted. On the contrary, both current and historical evidence suggests we may have at least as much reason to *start* with a more optimistic view of the ability of lay citizens to learn to navigate informational environments in which the number and range of informational sources available to them is rapidly increasing. Likewise, we should be open to the possibility that 21st century communications technologies are *improving* the epistemic well-being of democracy, and are likely to continue to do so in the future. If Progressives wish to make a compelling case to justify their deep pessimism about lay citizens or their pessimism about the epistemic impact of 21st century communications technologies on democracy, they need to keep in mind Criteria 1 given above, that our assessments of the epistemic well-being of democracy ought to be grounded in

rigorous and appropriate empirical evidence. Until then, and thankfully for the prospects of liberal democracy, we have just as much reason for optimism as pessimism.

Progressive Neglect of Criterion 2: Twenty-First Century Communications Technologies' Under-Theorized Relationship to Climate Denial, Anti-vaccers, and Newspaper Economics

Criterion 1 – making sure our claims about the epistemic impact of 21st century communications technologies (for example) are based on appropriate and sufficiently rigorous empirical research – is not the only requirement among those I provided above that is not met by Progressives who view 21st century communications technologies as a threat to democracy. Such commentators also often fail to fulfill Criterion 2, which is to *articulate the counterfactual* state of affairs against which the current epistemic state of democracy is being compared.

The lack of counterfactual reasoning in Progressives' pessimistic commentary about the epistemic impact of 21st century communications technologies on democracy can frequently be found in the commonly advanced Progressive proposition that the influx of misinformation whose spread is facilitated by the internet and social media has *caused increasing numbers of people to abandon or resist the scientific consensus* on matters of public importance. Perhaps the most widely discussed examples have to do with “anti-vaccers,” who resist the scientific consensus regarding the utility and safety of vaccines, and climate deniers, who resist the consensus of climate scientists on the reality and implications of, and human complicity in, climate change. To say that *increasing* numbers of people are abandoning or resisting the scientific consensus begs the question, “Compared to what?” If the answer is simply “compared to the way things ought to be,” that is probably true, but it is also insufficient grounds for finding 21st century communications technology guilty of harming the epistemic well-being of democracy. The normative standard according to which we ought to assess the present state of the epistemic well-being of democracy is not the best-possible state of affairs we can imagine in our minds but some alternative state of affairs we could *reasonably expect to exist in the absence of the internet and social media*. In this case, commentators who view the internet and/or social media as responsible for the sub-optimal influence of scientists and professional journalists on society need to make the case that, either (a) in some *past* era without the internet or social media or (b) in some *plausible alternative present* in which 21st century communications technologies

either never came about or took some different form, public acceptance of the scientific consensus *would be greater than it is at present in a way that does not make us worse off epistemically in other, more important ways*. Not only do commentators nearly always fail to articulate such a counterfactual (Criterion Two), however, but such failure often entails a simultaneous failure to ground their arguments in sound empirical evidence (Criterion One).

Let's look at these two issues in turn, beginning with climate change. The American public's intransigence when it comes to correctly identifying and accepting the scientific consensus on climate change is often held up as a paradigmatic example of the dangers of our informational environment, with commentators arguing that the overloaded information sphere makes it excessively difficult, in the first place, for the average citizen to figure out just what that consensus *is* and, beyond that, whether to *agree* with it or some alternative point-of-view to which they have easy access through the internet and social media. As William Minozzi puts it, “an increasingly competitive informational environment often provides more chances for jamming,”¹⁴² where by “jamming” he means impeding lay persons from recognizing and accepting the truth by sending out an alternative, untrue message.¹⁴³ If Minozzi's “jamming” theory of politics is true, we should expect the ever-increasing ease with which low-quality purveyors of information can publish contestations of scientific consensus to coincide with increasing public uncertainty about, or perhaps opposition to, that consensus. This is because, per Minozzi, in such an environment, “[a]dditional informational providers are empowered to cast an effective veto by disputing any information” they prefer not to gain public acceptance.¹⁴⁴ Minozzi's claim that simply cluttering up the informational environment with alternative claims equates to “casting an effective veto” only makes sense if he views lay citizens as fundamentally incapable of or unwilling to *make sound judgments* about what is true when presented with competing claims, and therefore reveals him to hold the same deep pessimism about lay citizens' judgment as other Progressives. After all, if citizens were capable of using their judgment to navigate competing claims with a fair amount of reliability, it would not necessarily matter *how many* alternative voices populated the informational environment. If citizens possessed good enough judgment, the mere *number* of voices cluttering the informational environment need not

¹⁴²Minozzi (2010), 311

¹⁴³*Ibid.*, 302

¹⁴⁴*Ibid.*, 311

necessarily be of much concern, since citizens possessed of sufficiently reliable powers of judgment should be able to ignore unreliable claimants and trust in reliable ones.

Were democratic citizens as unreliable judges as Minozzi's "jamming" theory, and Progressives' depictions of the public in general, imply them to be, then we should expect public trust in reliable sources like that of the "scientific community" to have decreased dramatically in the internet era. After all, as Progressives frequently point out, no era has remotely rivaled the Information Age in terms of the proliferation of competing voices it has allowed to clutter up the informational environment and the ease with which lay citizens have been enabled to access those voices. But despite this pessimism, and despite frequent insinuations by Progressives to the contrary, it is not at all clear that this is what has actually happened. Gallup data, for instance, shows that nearly 50% more Americans are *aware of* the scientific consensus about climate change now (71%) than in 1997 (48%). This increase coincides neatly with the time frame during which the internet became a household technology and a pervasive part of lay citizens' lives. The same study shows that on each of six measures related to global warming – whether it is human caused, whether most scientists believe it is occurring, whether its effects have already begun, whether it will pose a serious threat in the near future, whether they worry a great deal about it, and whether the news underestimates its seriousness – American attitudes between 2015 and the present have *fallen more in line with* the scientific point of view than was true from 2001-2014.¹⁴⁵ These trends, in turn, coincide with the proliferation of social media and its period of greatest integration into people's daily lives and routines.¹⁴⁶ In sum, at the same time that the internet, social media and mobile telecommunications technologies have experienced their greatest rates of expansion and services like Wordpress, Facebook, Twitter and Youtube have made it easier and cheaper for actors to flood the informational environment with "jamming" messages, Americans' *awareness of* and *agreement with* the scientific consensus on climate change have all *increased*.

Of course, strictly speaking, the proper counterfactual is not *whether* these attitudes have fallen more in line with the scientific consensus, but whether the *rate at which* they have tended to do so is lower in the informational environment served by the internet and social media

¹⁴⁵Jones and Saad (2017)

¹⁴⁶Ortiz-Ospina (2019)

compared to what would be the case otherwise. How might we estimate what the rate of scientific influence *would have been* increasing in society had the internet and social media not intervened? One obvious way of imagining this counterfactual would be to look at how readily and consistently public opinion fell in line with the scientific consensus in the era that *immediately preceded* the advent of the internet and social media, the broadcast era dominated by the legacy news media, which is frequently implied by Progressives to have been more friendly to science than the present.

Was that era in fact friendlier to science in this way? A look at the literature on public perceptions of climate change would suggest not. On the contrary, the informational environment that immediately preceded the heyday of the internet and social media does not appear to have been especially amenable to even making citizens *aware of* the scientific consensus on climate change, let alone getting them to *agree with* it. In an article published in 1992, Sheldon Ungar reported that before 1988, scientists had “tried for more than two decades to attract attention to global warming” but had been largely ignored by the public. According to Ungar, citizens' attention was brought to the issue of global warming not by the timely reporting of professional journalists or the effective curation of information by editorial boards and television news producers, but instead by the extreme heat wave felt nationwide during the “greenhouse summer of '88,”¹⁴⁷ a verdict that was also endorsed six years later by Richard J. Bord and colleagues. The latter group, however, also found that after the heat wave dissipated and stopped making climate salient to citizens' personal lives, climate change receded dramatically in importance in the eyes of the public. By 1997, only 43% of Americans agreed with the characterization of global warming was even a “somewhat important” issue.¹⁴⁸ Meanwhile, both a Cambridge poll in 1994 and a Gallup poll in 1997 found that a majority of Americans were not even aware *that* a majority of scientists believed that global warming was occurring *at all*.¹⁴⁹ That is to say that, after attempting to raise awareness of the problem for over *three decades* in the pre-internet and pre-social media informational environment, the “scientific community” had by the end of the broadcasting era dominated by the legacy news media not even managed to apprise the majority of the public of *what* it thought about global warming, let alone persuade the public that *they*

147Ungar (1992), 483

148Bord et al (1998), 77

149Nisbet and Myers (2007), 452

ought to be worried, too. By contrast, polls now consistently find that a substantial majority of the public is aware of the scientific consensus on climate change, agrees that global warming is at least partially caused by human activities, and even that its effects are already being felt.¹⁵⁰

Now, it could be that these increases would have been more pronounced if the internet and social media had not been around or if they had arisen in some different form or been managed in some different way. However, that is a counterfactual case that has to be carefully prosecuted by those condemning the epistemic impact of these technologies when it comes to societal awareness of and agreement with the scientific consensus on climate change; it is not something we can just assume. Yet, despite the voluminous commentary that attributes a causal role to the internet and social media in maintaining public ignorance of and resistance to the scientific consensus on climate change, I have yet to encounter a single instance in which such a counterfactual case for what *would otherwise have been true* of the present, were it not for the internet and/or social media, has been constructed. Most of that commentary does not even make the effort to compare the difficulty *currently* faced by climate scientists and their allies in getting the public's attention and buy-in with what was true *before* the advent of the internet and social media. In other words, none of it fulfills Criteria 2, the requirement to articulate the counterfactual against which the current epistemic state of democracy is being compared. The absence of such counterfactual reasoning critically undermines the oft-forwarded claim that the internet and social media are somehow causally responsible for continued public opposition to the climate science consensus. Indeed, if anything, the evidence cited above suggests that the case may be just the opposite – that an informational environment that includes widespread internet and social media use actually *facilitates* awareness and agreement with that consensus.

Of course, climate change is not the only issue Progressives frequently cite as a case in which 21st century technologies have undermined lay citizens' awareness of or trust in the scientific consensus. Much has also been made of a few highly publicized groups and movements who have abandoned the scientific consensus on certain matters in favor of outlandish alternatives, such as the “anti-vax” movement whose members view vaccines as dangerous and/or ineffective, with many again implicating the internet and social media as facilitators of these beliefs. For instance, Larson and colleagues write that anti-vaccination

¹⁵⁰Jones and Saad (2019)

movements have attained “new levels of global reach and influence” due to their having been “empowered by the internet and social networking capacities.”¹⁵¹ Is the case for pessimism about the epistemic impact of 21st century communications technologies on democracy any stronger with regard to public attitudes toward vaccination than it is for climate change? I do not think it is. The Center for Disease Control's (CDC) vaccination coverage report for 2017 (the latest year for which data are available), shows that since 2012 “Overall vaccination coverage among young children [has] remained high and stable in the United States,” and a look at the data shows that, if anything, vaccine coverage has actually *increased* over that duration, which, again, also happens to cover the time period during which social media use experienced its greatest-ever rates of growth and overall usage.¹⁵² There *are* isolated communities ²where comparatively large numbers of parents refuse to vaccinate their children, resulting in measles outbreaks that have gained national media attention, specifically, in King County, Washington¹⁵³ and several orthodox Jewish communities in Brooklyn.¹⁵⁴ These cases are usually the ones that are focused on by Progressives worried about the direction of society with regard to vaccination in the 21st century. However, anyone acquainted with statistical reasoning should know that the existence of a few communities with abnormally high numbers of non-vaccinated individuals by no means proves an *overall* increase in the influence of anti-scientific thinking in the public. On the contrary, we should *expect* there to be a few such communities as a simple matter of probability. Moreover, the geographic isolation of the outbreaks in King County and orthodox Jewish communities of Brooklyn calls into question the role played in the development of such vaccine-resistant communities by the internet and social media. If the internet and social media were the major reasons why the citizens involved in these outbreaks resisted vaccination, and not other social factors such as parental upbringing or the religious or cultural environment of these communities, why should their effects be so geographically isolated? The consistently high rate of vaccine coverage in the digital age, the scarcity of outbreaks like those in King County and the Orthodox Jewish community of Brooklyn, and the geographic isolation of those outbreaks, all

¹⁵¹Larson et al (2011), 526

¹⁵²Hill et al (2018), 1126

¹⁵³Washington Department of Health (2019)

¹⁵⁴Belluz (2019)

call into question the extent to which 21st century communications technologies may be said to be causing some sort of large-scale resistance to the scientific consensus on vaccination.

As was the case with the claim that the internet and social media are friendly to unreliable sources of information and unfriendly to reliable ones dealt with in the previous sub-section, these contemporary considerations are buttressed by historical considerations. The existence of isolated anti-vaccination communities is not some revolutionary contemporary development unique to the digital age. On the contrary, such communities have frequently existed even in developed countries ever since vaccinations became widespread in the developed world.¹⁵⁵ Not only does the existence of isolated under-vaccinated *communities* precede the internet. So do anti-vaccination *movements*. Ironically, their main cause appears *not* to have been the sudden proliferation of anti-vaccination materials by ingoramuses, “extremists” and manipulators suddenly given a voice by modern technology, but instead the exact opposite – that is, arguments against vaccination made *by scientists and other sources considered to be respectable epistemic authorities*. An international study compiled by Gangarosa and colleagues identified numerous anti-vaccination movements around the world between 1985-1995. In every one of the eight case-studies in which anti-vaccination movements resulted in lower vaccine coverage among an appreciable number of citizens – i.e. in every case in which a demonstrable, widespread effect on the *behavioral influence* of the movement could be detected – the authors identified either the dissemination of anti-vaccination sentiment by a prominent member of the medical community, disagreement between authoritative institutions on the effectiveness and side-effects of vaccines, or the “attitudes, knowledge, and practices of physician providers” as the primary cause. As Gangarosa and colleagues point out, these anti-vaccination leaders were given a voice by the legacy news media, book publishers, and even scientific journals, of their day,¹⁵⁶ which undermines the implication that the pre-internet informational environment was especially good at filtering out unreliable sources of information and augmenting the voice of reliable ones. Just as was the case for climate change, then, is it neither clear *that* opposition to the scientific

¹⁵⁵For instance, Japanese authorities seeking to study whether pertussis' effects in the modern age had declined so much as to justify the discontinuation of the pertussis vaccine regime – a common hypothesis in the 1970s-1980s (Gangarosa et al, 1998) – found that the instances of pertussis it discovered were isolated to “24 cases in Kanagawa prefecture, where small epidemics persisted in Miura area since the preceding year” (Kanai 1980, 112). The recent polio “outbreak” in Nigeria likewise consisted of “several cases of poliovirus in the state of Borno in northeastern Nigeria” (Center for Disease Control and Prevention 2019).

¹⁵⁶Gangarosa et al (1998), 358-9

consensus regarding vaccines has increased *at all* in the information age nor that 21st century communications technologies would necessarily be responsible for any change that *did* occur. Again, commentators may have some plausible counterfactual world in mind when they argue or imply that a world without, or with very different versions of, the internet and social media would be one in which anti-vaccination movements would have a harder time forming and flourishing than the one that existed immediately before the advent of these technologies. However, if that is the case, those commentators need to articulate and defend that counterfactual scenario.

Both in the case of climate denial and in the case of opposition to vaccines, then, not only is it not clear that the internet and social media are hampering the Progressive aim that citizens become aware of and demonstrate a willingness to believe the scientific consensus, but the case may even be the *opposite* – the informational environment dominated by these communications technologies may even be *conducive to* these aims, even in its current warts-and-all form, while the pre-internet informational environment dominated by publications institutions and the legacy news media may not have been very good at encouraging public awareness and acceptance of the scientific consensus. At a minimum, this means that we should resist the conclusions of those who explicitly or implicitly blame the internet and social media for harming the epistemic well-being of democracy by undermining these knowledge-producers and the industries that support them until Progressives lay out a convincing counterfactual case.

Again, I will close this sub-section by noting the latent pessimism about the judgmental capacities of lay citizens that infuses Progressives' expectations and interpretations of 21st century communications technologies' effects on democracy. Both Minozzi's "jamming" theory of politics and Progressives' depictions of the likely impact of 21st century communications technologies on democracy *presuppose* a deeply pessimistic view about the judgmental capacities of lay citizens by implying the the *mere* entry of more and more alternative voices into the informational environment is bound to be detrimental to the public's becoming aware of and accepting the scientific consensus on climate change and vaccination, respectively. The *mere* entry of new voices into lay citizens' informational environment should only be expected to lead citizens to abandon reliable beliefs and embrace unreliable ones if those citizens' judgment cannot be depended on to make good choices about what to believe *by the lights of their own*

judgment. Progressives' claims that 21st century communications technologies are an impediment to lay citizens' becoming more aware and accepting of the scientific consensus on climate change and vaccination stem largely from their having *presupposed* that citizens' judgment is deficient in this way. Without this presupposition, however, the case for Progressive pessimism about the epistemic impact of 21st century communications technologies regarding public awareness and acceptance of the scientific consensus on climate change and vaccination becomes very tenuous, if anything running against the grain of both current empirical evidence and historical trends.

Neglect of Criterion 3: Our Two Epistemic Goals – Improving Expertise and Expanding Citizens' Awareness of Politically Relevant Knowledge Beyond the Contents of Legacy News

So far in this chapter, I have joined Progressives in endorsing the idea that one important part of maintaining the epistemic well-being of democracy is for citizens to become aware of and be willing to place their trust in the claims of scientists and professional journalists. In line with the Progressive Ideal articulated in Section 2, it is common for commentators to argue or strongly imply that citizens ought to place their trust in these and various other experts, i.e., persons endowed with credentials or occupying a social position that signifies some kind of epistemic authority. In this sub-section, I will highlight the limitations and dangers inherent in endorsing widespread public trust in those endowed with such credentials or occupying positions that imply that kind of epistemic authority.

I will start off by noting that there are numerous examples of individuals and institutions considered experts by both lay citizens and the intelligentsia of their time who turned out to be purveying false information. I doubt anyone would contest that this is true of “officials,” broadly construed. Many would also agree that this has been true of at least some professional journalists. However, I expect more to take umbrage at the idea that this might be true of scientists. While I agree with the general sentiment that the “scientific community” has historically been one of the most reliable sources of knowledge in the contemporary – or probably *any* – era, even it is not immune from criticism on epistemic grounds. In fact, we have actually encountered one example of a way members of the scientific community have contributed to the spread of faulty beliefs, which is by fueling opposition to vaccines. One of the most commonly cited arguments forwarded by anti-vaccers is the claim that the MMR (measles,

mumps and rubella) vaccine causes autism. That claim gained its original momentum on the basis of an article published in the prestigious medical journal *The Lancet*. While *The Lancet* eventually published a retraction of that article, it is surely no coincidence that one of the most popular arguments forwarded by anti-vaccers is a claim that originated in the scientific community. Such claims, after all, are nearly universally considered by the public to be more reliable, at least on matters like virology, than the claims of lay persons, amateurs, and members of intellectual communities outside of science, such as philosophical societies. Nor is this danger necessarily averted by recommending that citizens place their trust not in the claims of individual scientists but instead the oft-invoked “scientific consensus.” Such “consensuses” often overlook the extent of disagreement that usually exists *within* the scientific community, even on matters commonly thought to be “settled.”¹⁵⁷ But even when the word “consensus” is appropriate, it is no guarantor of correctness.¹⁵⁸ On the contrary, as the earlier discussion of gastroenterologists' opposition to any research attributing peptic ulcers to bacteria demonstrates, it can be a sign of dogmatism and closed-mindedness.¹⁵⁹

While I take both the *Lancet* article and the gastroenterologist community's opposition to the bacterial hypothesis to be the exception rather than the rule, I also believe that they demonstrate the persistent dangers entailed in sanctioning too much trust in those who bear the trappings of epistemic authority. As I will stress more forcefully in the later chapters of this dissertation, I also think it is reasonable to expect that experts and expert communities will become *less* reliable the more they realize they can count on members of the public regularly deferring to their expertise even when those experts' claims strongly diverge from their own judgment, a disposition that seems to be very close to what Progressives advocate. This is because such quiescence is an open invitation either to abuse or at least to complacency, and we

¹⁵⁷See, e.g., Collins and Pinch (2007)

¹⁵⁸Another contemporary example can be found in the field of nutrition, whose views on the role of dietary fat in obesity and heart health have changed dramatically in the past two decades. I cannot help being reminded of my noticeably trim and energetic grandparents' stodgy refusal to stop eating bacon and eggs, reasoning that they had eaten them their whole lives and were perfectly well placed to tell what impact they were having on their bodies.

¹⁵⁹Another series of examples of why limits ought to be placed on citizens' willingness to trust scientists can be found in the shocking revelations that have been made in recent years about sexual abuse committed by sports physicians in various universities. In the most famous such case, that of former Michigan State and Team USA gymnastics physician Larry Nassar, his molestation of underage girls was “explained at the time as treatment” (Tucker 2017). In other words, Nassar was invoking his epistemic authority as a scientist to take advantage of lay citizens in ways that harmed the physical and psychological health of children.

have every reason to believe that such an invitation is likely to be taken up sooner or later. These dangers of abuse and complacency are encouraged by commentators who place excessive emphasis on the democratic value of deferring to epistemic authorities and insistently characterize those who voice opinions that differ from the views endorsed by such authorities as a threat to democracy's epistemic well-being.

I believe the tendency to adopt this one-sided approach and to overlook the risks entailed in too widespread a readiness to defer to the claims of experts often stems from researchers' failure to fulfill Criterion 3 outlined above, which is to take into account *both* our epistemic aims – the aim of *discovering truths* as well as the aim of *avoiding mistaken beliefs*. When commentators portray small pockets of citizens who subscribe to eccentric beliefs – such as the belief that the world is flat, that the holocaust never happened, or that vaccines either do not work or cause even worse diseases – as threats to the epistemic well-being of democracy, they seem to be embracing the notion that the more people who believe what is currently taken to be the truth by the relevant epistemic authorities – say, for instance, scientists – the better off democracy is epistemically. The logic that lies behind this notion seems to go as follows:

1. **Premise 1:** At least when it comes to matters that lie outside common experience, experts are better at finding and disseminating the truth than lay persons.
2. **Sub-Conclusion:** When it comes to such matters, lay citizens maximize their odds of believing the truth by accepting the claims of the relevant experts, save in exceptional circumstances.
3. **Premise 2:** Democracy's epistemic well-being is improved the more each individual citizen maximizes their odds of believing the truth about the greatest number of things.
4. **Conclusion:** Democracy's epistemic well-being is improved to the degree that more and more citizens are willing to accept the claims of the relevant experts (save in exceptional circumstances).

As intuitive as this logic may seem, I believe it is mistaken. The key mistake lies in Premise 2, which views making decisions about what to believe as a *one-off, individualistic* process that is properly approached in a *risk-averse* manner. Now, it is true that in many

situations it makes sense to follow this approach. Specifically, it makes sense when (1) an *individual* is (2) faced with a *single decision* where (3) getting it wrong carries significant *consequences* and (4) there seems to be a *strong and apparently well-founded* consensus among the relevant expert community that there is a known solution to the problem. But if any of these variables change – if the relevant decision-making entity is a *collective*, for example, or if an individual or a collective is faced with a *series of decisions* over time where the risk of getting a single one wrong is not enormous – the appropriateness of Premise 2 is thrown into doubt. For example, as I will elaborate at further length in Chapter Three, studies by social epistemologists have consistently shown that when it comes to finding optimal solutions to problems, knowledge-seeking collectives – the scientific community is the paradigm example – in which individual researchers or research groups' pursuit of the optimal known solution to a problem at any particular time is *limited*, due either to ignorance of that solution or a dogmatic refusal to immediately adopt it, *regularly and dramatically outperform* knowledge-seeking collectives in which the most-likely solution at a given time is widely accepted and acted upon by each individual researcher or research group.¹⁶⁰ The basic reason is that even the best of our present knowledge is limited and fallible, which means that (a) the expert consensus in any given instance might be wrong and (b) even if it is “right” in the sense that it will help us solve whatever problem is before us better than any other as-yet discovered solution, it might not be the *best possible* solution to that problem. If, as is nearly always the case, either (a) or (b) is true, then the relevant decision-making body – in our case, democracy – stands to epistemically benefit from the persistent willingness of at least *some* individuals or groups to committedly, even obstinately, dissent from the current epistemic authorities' consensus, *even when* it seems obvious to everyone that that consensus is correct. This is not to say that there are no benefits to the widespread adoption of common beliefs by the public. On the contrary, as John Stuart Mill put it in *On Liberty*, “the well-being of mankind may almost be measured by the number and gravity of the truths which have reached the point of being uncontested.” However, Mill is quick to point out in the same passage that this does *not* imply that universal acceptance of (what we currently take to be) the most strongly supported beliefs is strictly beneficial, writing that

¹⁶⁰For examples in social epistemology, see Weisberg and Muldoon (2009) and Zollman (2010). For applications of these and similar concepts to politics, including references to many more studies, see Page (2007) and Landemore (2012).

although the “gradual narrowing of the bounds of diversity of opinion...is at once inevitable and indispensable, we are not therefore obliged to conclude that all its consequences must be beneficial.”¹⁶¹ Instead, it comes with the drawbacks of diminishing the likelihood that we will discover the errors in our current beliefs, losing sight of the strongest grounds for holding those beliefs, and discovering which aspects of those beliefs we currently hold that are mostly justified are built on shaky foundations. I take this line of reasoning to be part of the reason James warned his scholarly contemporaries in the early 20th century of the danger of neglecting our epistemic imperative to *believe truth* in our determination to *avoid error*. Were *every* individual or group to *always* make the risk-averse move of attempting to maximize their chances of avoiding error by embracing the current expert consensus, we as a society would inhibit our ability to discover new and unanticipated truths – truths like the existence of a new and better solution to a current problem, or perhaps corruption or complacency among a body of epistemic authorities we previously thought to be trustworthy.

Nothing in what I have said should be taken to suggest that dissenting from expert opinion should be the societal *norm* or that such dissent is in every *particular instance* an epistemically responsible move on the part of the dissenting individual or group. Nor does it imply that we ought not to worry about making errors at all. As James said, discovering new truths and avoiding errors are *both* epistemic commandments. Whether we as a democratic society are striking the appropriate balance between these two goals is bound to be the subject of debate in any given instance. But whatever balance we settle on, it must be a *balance*, meaning that we oughtn't *ignore* one of these commandments out of a fixation on the other. The point is that the *mere existence* or even the *moderate growth* of what we take to be wacky, ignorant, blindly dogmatic or otherwise epistemically irresponsible groups is not necessarily a bad sign for the epistemic well-being of democracy, as it is commonly argued or insinuated to be.¹⁶² In fact, it may be just the opposite – a sign that our society is set up in a way that increases our odds of discovering new and unanticipated truths. I shall stress the implications of this argument at greater length in Chapter Three.

¹⁶¹Mill (2007 [1859], 101)

¹⁶²See, e.g., Sunstein and Vermeule (2009), Sunstein (2016), Lewandowsky et al (2017)

For now, however, I will move on to another point. Even were the line of reasoning above not sufficient to undermine the argument that the epistemic well-being of democracy is better off to the extent that lay citizens are willing to put their trust in those currently taken to be the experts in the relevant epistemic domain, there are other ways in which the distinction between discovering new truths and avoiding errors can help us to see the epistemic value of the internet and social media in ways that a myopic focus on error avoidance cannot. One way of phrasing the point, applied to the topic of democracy's epistemic well-being, would be to say that democracy's epistemic well-being depends not only on the *percent* of truths that fill up its informational space but also the *number* of truths that space allows citizens to become aware of. Consider professional journalism of the sort viewed by Progressives as integral to the epistemic well-being of democracy. That profession's democratic importance stems largely from its role in keeping citizens informed about significant events, societal trends, the actions and words of politicians, and so forth. Now, there can only be so many professional journalists. At its apex, before the total number of newsroom staff began to plummet around the turn of the millennium, the total number of persons employed by the legacy news media was estimated by Pew to be around 64,000.¹⁶³ This means that in the pre-internet era, the news items that the public at large could come to be made aware of were largely limited to those events that reached the networks built up by these employees and those of wire services like the AP and Reuters. Extensive as they may have been, such networks were nowhere near comprehensive or reliable enough to effectively monitor *all* of the world or even *all* of the United States. Instead, reporters did what we all do when faced with problems that require more data than we can hope to gather ourselves. They used shortcuts. Reporters focused their networking on sources of information that seemed especially likely to generate what came to be generally accepted as “the news,” such as local crimes and the words of prominent politicians. During this time, journalists were often blasted by critics for relying so heavily on such a homogenous and predictable staple of information sources and information-gathering routines, which, critics claimed, produced a skewed picture of the world and make them easily manipulable by politicians and PR firms.¹⁶⁴ I believe these are real and important limitations to professional journalism as it was practiced in the 20th century, but I

¹⁶³Pew State of the News Report (2004)

¹⁶⁴See, e.g., Cook (1998), Sparrow (1999), Bennett (2016)

also believe that it is hard to imagine what the legacy news media could realistically have done otherwise given the enormity of the epistemic task before them. There is simply no way the legacy news as an institution could have actually apprised itself of and informed the public about *everything* of political relevance to democratic citizens. It is very doubtful it could even have managed this for every *important* eventuality of political relevance. Their choice to cover the news in the way they did was driven by the sheer magnitude of the task set before them. Even, then, if journalists and editors were at times guilty of following them too uncritically, the news' adoption of norms, routines and beats was an understandable and perfectly defensible response to the complexity of the task of monitoring an increasingly large and complex society.

The basic obstacle professional journalists sought to overcome with norms, routines and beats was the enormity and complexity of society. Sixty-four thousand journalists could only be in so many places at once. No matter how enterprising their reporters and how seriously they took their democratic duty, their epistemic resources were limited by basic constraints of time and space. Now, contrast this meager set of epistemic resources against the collective epistemic resources of *democratic citizens* with access to Wordpress, Facebook, Twitter, Youtube and the rest of the internet's resources for cheaply and rapidly disseminating and accessing information. The potential for the billions of users of these 21st century communications technologies to monitor and report on events and phenomena of relevance to democratic citizens vastly outstrips the potential of the legacy news media, even at its zenith, to accomplish that same task. Indeed, the collective awareness of the billions of citizens who now have access to social media and other 21st century communications technologies is by now so extensive as to represent a potential awareness of the vast majority of what happens in the social world. As such, social media and the blogosphere represent an enormous increase in our collective potential to become apprised of a far greater quantity and variety of knowledge of precisely the sort professional journalists sought so long to monitor for the public – that is, knowledge of significant events, societal trends, the actions and words of politicians, and so forth. And, indeed, no one seems more aware of this than the legacy news media itself, which increasingly relies *on* social media and the blogosphere both for content and for “leads” to follow up on. Even if we could *always* count on professional journalists to get *perfect* knowledge about whatever matter they set their efforts to studying, and even if an informational environment without the internet or social media would be better at

communicating them accurately to the public – neither of which assumption I take to be clearly justified – the sheer *magnitude* of the increase in the volume of information made available to citizens by the internet and social media may *still* be so great that the tradeoff would be worth it. That is, even if the internet era's informational environment is populated, as Progressives frequently portray it to be, by both a greater *number* and *percentage* of unreliable claims, it could be worth putting up with those because of the massive increase in the sheer *quantity* of *reliable* information to which those same technologies give citizens access. In order for Progressives to justify their claims that 21st century communications technologies are harming democracy epistemically, they need to take into account *both* our epistemic commandments, and not just fixate on citizens' need to avoid adopting unreliable beliefs.

Once again, it is not clear that the Progressive approach is justified by either current empirical evidence or historical trends. While some researchers have implied that social media environments such as those of Facebook¹⁶⁵ and Twitter¹⁶⁶ are as full of fake/false news as truth, many other studies strongly suggest that fake news is in fact very rarely read or shared.¹⁶⁷ As for misinformation conceived more broadly, very little research has been done to try to actually quantify the proportion of claims disseminated online are true versus false. The single empirical study I know of in this vein, a study published by Vosoughi and colleagues in *Science* that suggests that falsehood spreads “significantly farther, faster, deeper, and more broadly than truth in all categories of information,”¹⁶⁸ is based entirely on claims that were “fact-checked by six independent fact-checking organizations.”¹⁶⁹ It should be obvious that such organizations are likely to focus their attention primarily on claims some group of people has seen fit to contest and to ignore claims that no one views as worth contesting, perhaps even on claims that can be refuted, since such refutations demonstrate the expertise of the individuals and organizations doing the “fact-checking,” and so the results are very likely to overestimate the presence of false/unreliable information versus true/reliable information on Twitter. The empirical case supporting Progressive pessimism at present, then, is as yet unconvincing. Likewise, the same arguments made above about the historical relationship between an increase in the number of

¹⁶⁵Silverman (2016)

¹⁶⁶Vosoughi et al (2018)

¹⁶⁷E.g. Fletcher et al (2017), Guess et al (2017), Nyhan (2018)

¹⁶⁸Vosoughi et al (2018), 1147

¹⁶⁹Ibid., 1146

sources of information easily available to the public and the accuracy of the beliefs subscribed to by that public undermine Progressives' implicit assumption that we ought to be more worried about the potential for people to adopt mistaken beliefs in such an environment than optimistic about their potential to discover new truths. It is far from clear that the beliefs of lay citizens have been more harmed by their access to false/unreliable information than they have benefited from increased access to true/reliable information.

As before, I will close by noting the thick vein of pessimism about the judgmental capacities of lay citizens implicit in Progressives' approach to the epistemic impact of 21st century technologies on democracy. Progressives' fixation on *misinformation* as the primary type of information relevant to assessing the epistemic impact of 21st century technologies on democracy reinforces the suspicion that they enter into their consideration of that topic with a view of the judgmental capacities of lay citizens that is *already* deeply pessimistic. After all, if they were *optimistic* about lay citizens' judgment, it seems like they would pay at least as much attention to citizens' greatly improved potential to *discover new (to them) truths* as to the danger of their believing falsehoods. Even if Progressives were not optimistic but merely *neutral* about the judgmental capacities of lay citizens, one would expect them to see it necessary to try to engage in a balanced consideration of the epistemic pros and cons of lay citizens having access to so many more, and more diverse, sources of information instead of fixating almost entirely on the increased presence of *misinformation* in the informational environment and the corresponding increase in lay citizens' opportunities to epistemically err. As before, the case for Progressives' pessimistic *conclusions* about the epistemic plight of democracy in the Information Age seems to be under-supported by current evidence and historical trends and to depend critically upon Progressives' pessimistic *presumptions* about the judgmental capacities of lay citizens.

Neglect of Criterion Four: The dangers of generalizing about “The Internet” and “Social Media”

As I mentioned in the discussion of Criterion One, it is common for Progressives to write as if digital technologies are unfriendly to both scientists and professional journalists while being friendly to less reliable ones. When they do so, they often use sweeping terms to generalize about

the *entirety* of “the internet” or “social media.” Journalism scholar Anya Schiffrin, for instance, characterizes “the internet” as “a forum for distributing information that does not adhere to typical standards of truth, scientific inquiry, and evidence-based news and information.”¹⁷⁰ While in this passage it is the internet as a *whole* that draws Schiffrin's disdain, more widespread is the view that *social media* in particular is the enemy of scientists, professional journalists, and other reliable sources of information. The British House of Commons Report on Disinformation and “fake news,” for example, warns that the consequences of people “increasingly finding out about what is happening...through *social media*, rather than through more traditional forms of communication” like television and newspapers could be “devastating” to democracy.¹⁷¹ Just as Schiffrin did with “the internet,” the House Committee writes as if making a *general* assessment of the reliability of the contents of “social media” is a sensible and useful way to discuss its relationship to the epistemic well-being of democracy. This tendency to generalize – and usually to generalize negatively – about “the internet” and “social media” is also common in public opinion research, mainstream political commentary and indeed among the public at large. In the wake of the 2016 presidential election, in which talk of fake news and Russian activity on social media was rampant, this pejorative view of social media as especially friendly to unreliable sources and unfriendly to reliable ones seems to have gained purchase in both America and in Europe public, with a recent Pew poll showing that “a majority (57%)” of Americans “say they expect the news they see on social media to be largely inaccurate”¹⁷² and a poll conducted by the YouGove-Cambridge Globalism Project revealing that in seven out of eight European countries more than half of respondents answered “not much” or “not at all” to the question, “How much, if at all, do you trust information from social media?”¹⁷³ This question's form implies to participants and readers that it makes sense to talk about the trustworthiness of the information citizens encounter on social media *in general*. As we will see in later chapters of this dissertation, this perspective often motivates Progressives to encourage democratic citizens to adopt a default disposition of skepticism toward the information passed on to them by their peers via the internet and social media, whereas they advocate a less skeptical disposition when it comes to the

170Schiffrin (2017), 123

171House Report (2018), 4, my emphasis

172Matsa and Shearer (2018)

173The single exception was Poland, clocking in at around 45% (Smith 2019).

information passed on to them by the “scientific community,” the legacy news media, and others Progressives consider to be legitimate epistemic authorities.

This habit of generalizing about “the internet” and “social media” as a whole is highly misleading and is a major contributor to commentators' tendency to fall short of Criterion Four listed above, which is that verdicts about 21st century communications technologies' epistemic impact on democracy need to *account for causal complexity*. By lumping together all the contents of the extraordinarily complex communications technologies like “the internet” and “social media” into a single category and attaching a single descriptor such as “unreliable” as a sort of summary of their “average” epistemic quality – an “average” that, as demonstrated in the previous sub-section, is not based on any kind of rigorous analysis but instead on the commentator's *impressions* – Progressives risk adopting conclusions about the epistemic impact of 21st century communications technologies on democracy that are likely to be not just misguided but *actively misleading*. This is true for the same reason that basing one's analysis on an extremely high-variance data set solely on the regression curve is likely to mislead analysts looking for meaningful trends. In both cases, it is the differences *between* given data points – in our case, individual web domains/pages or individual social media accounts/posts – that is likely to be probative when it comes to understanding the important aspects of the subject matter, rather than a simple averaging of the data as a *whole*. When a dataset has very high variance, it often means that the causal relationship between the outcome of interest (dependent variable) and the phenomenon proposed by the researcher's model as a contributor to that outcome (explanatory variable) is weak, and that therefore researchers concerned with explaining or changing the outcome ought to look for other phenomena. I believe this is exactly the approach we ought to take in studying the epistemic well-being of democracy and the effects of 21st century communications technologies on it because the relationship between 21st century communications technologies is *causally complex*, by which I mean both that *different aspects of* 21st century communications technologies have *different effects on* the epistemic well-being of democracy and that *the same aspect of* 21st century communications technologies can have *multiple, different effects* on the epistemic well-being of democracy. In this sub-section, I hope to demonstrate that the Progressive tendency to over-generalize about the epistemic quality of 21st century communications technologies and to fail to account for causal complexity in their

analyses of those technologies' epistemic impact on democracy contributes to their excessive – or at least premature – deep pessimism about lay citizens and the direction of democracy.

Consider for a moment what is implied when one issues a proclamation regarding, or asks a citizen to offer their opinion about, the quality of the information found on “the internet” or “social media,” writ large. This approach seems to imply that there is enough uniformity between the sorts of claims one might expect to find on *the internet* and/or *social media* to make such generalizations sensible. In the examples given above, both the House of Commons committee and Pew seem to back into this way of talking about the internet and social media by comparing their reliability to professional news outlets. But does it make sense to generalize about the “average” reliability of the claims published via *communications technologies* like the internet and social media in the same way that it makes sense to generalize about the reliability of the claims published by *knowledge-disseminating institutions*, such as the “scientific community” or the legacy news media? I think it obviously does not. The processes by which a given article or claim may come to be published on the internet and social media differ dramatically from those required to gain publication via knowledge-disseminating institutions like the legacy news or a scientific journal. Articles that eventually appear in the *New York Times* or *Science* first go through a well-established, uniform and rigorous *editorial process*, which includes each article being subjected at numerous points to the judgment of highly trained peers. Posts that appear on the internet or a social media user's “wall” or “feed,” on the other hand, are subject to nothing remotely resembling the same degree of uniform procedures. Instead, some of these contents undergo some sort of rigorous editorial process prior to publication, others are subject to no process at all, and others go through some process that lies between these extremes. What is more, the *type* and *origins* of work *submitted* for public viewing on the internet and/or social media varies enormously, whereas work submitted to the *New York Times* or *Science* is highly homogeneous in both respects – its subject matter is constrained and its writer is typically someone acculturated into the standards and norms of professional journalism or some scientific field. The uniform adoption of this content and these standards and norms imposes a high degree of homogeneity on the works produced by professional journalists and scientists even *before* they go through the further-homogenizing editorial systems of these respective institutions. Again, this homogeneity of occupation and acculturation is utterly unlike what characterizes the hodge-

podge of pages and posts that make their way onto the internet. Given the homogeneity of the work *submitted to the New York Times and Science* as well as the uniformity and rigor of the *editorial system* through which such submissions are subsequently *curated*, it makes sense to infer that there will be a certain degree of commonality among the publications produced by these and other knowledge-disseminating institutions characterized by similar levels of homogeneity on both fronts. However, the exact opposite is true for the internet and social media. Given the extreme variety of authors whose claims appear on the internet and social media, and given that there is nothing remotely resembling a uniform, rigorous editorial process their posts must undergo in order to get published, it makes no sense to generalize about the reliability of the contents found on “the internet” or “social media” in general terms.

Generalizing in sweeping terms is not just a poor way of discussing the *reliability* of the information found on the internet and social media; it is also very likely to be a poor way of discussing the *belief-formation process* undergone by citizens encountering that information. It is extremely unlikely that the decisive factor in the minds of most citizens assessing the credibility of a given claim is whether that claim was found *on the internet* or *on social media*. There has been a considerable amount of empirical research into online credibility evaluation. That research strongly suggests that rather than focusing on whether claims appear *on* the internet or social media, people look to a highly variegated array of cues when deciding which claims to believe.¹⁷⁴ These cues are not only numerous and diverse in kind, but *which* cues an individual relies on *itself* varies depending on a number of factors such, for instance, as prior familiarity with the subject matter and overall education level;¹⁷⁵ whether a user encountered a claim by deliberate search or happenstance;¹⁷⁶ and whether they happen to trust the official sources available on the topic at hand.^{177,178} Even if it did make sense to generalize about the reliability of

¹⁷⁴Metzger et al (2010)

¹⁷⁵ Lucassen et al (2013)

¹⁷⁶Klawitter and Hargittai (2018)

¹⁷⁷Lederman et al (2014)

¹⁷⁸If the sort of crude heuristic suggested by commentators who summarize the reliability of “the internet” or “social media” in sweeping terms is employed by anyone, it is likely to be those who are largely technologically illiterate, such as older citizens who did not grow up using digital technologies. And indeed there is some evidence that older citizens are worse at navigating truth-claims on the internet than those who grew up in more digitized eras. Guess and colleagues found that when it came to sharing fake news during the 2016 election, the variable of age was highly predictive while “none of the other demographics variables in our model – sex, race, education, and income – have anywhere close to a robust predictive effect” (Guess et al 2019, 2). This may well be because they rely on outdated heuristics that are inappropriate to the internet era, such as

information found on the internet or social media as a whole in the first place, then, that would still not necessarily tell us anything about the epistemic well-being of democracy because it would tell us nothing about which particular pieces of the information available in that environment influenced the beliefs of citizens.

The tendency to compare the internet and social media to knowledge-disseminating institutions like legacy news networks also relates to another way in which commentators on the relationship between 21st century communications technologies and the epistemic well-being of democracy tend to overlook complications that arise from the complexity of the subject matter. This is commentators' tendency to treat communications technologies like the internet and social media as *alternatives to* reliable knowledge-disseminating institutions like the scientific and professional journalistic communities. This tendency is evident in the quotes cited above by Schifffrin and the House of Commons Reports and is far from uncommon in mainstream commentary about the internet and social media's impact on democracy.¹⁷⁹ But of course, neither professional journalism nor science is *distinct from* or should be treated as *mutually exclusive with* either the internet or social media. On the contrary, the work of both professional journalists and scientists appears *on* the internet and social media with great regularity. In fact, it seems a safe bet to say that the internet in its short lifespan has already been responsible for a greater and more extensive distribution of scientific knowledge than any technology that has ever existed. Meanwhile, every major legacy news company has its own website through which it distributes news. In fact, the digital audience of even the most well-established newspapers now vastly outstrips its paper audience. The *New York Times*, for example, in 2018 reported having less than a million paper subscribers, while its digital subscribers numbered over three million, a proportion that undoubtedly vastly underestimates the total digital distribution of the *Times'* stories since visitors are able to access several articles per month without a subscription and anyone can see the headline and summary of the articles linked by their friends on social media. As the latter observation suggests, even social media, much maligned for the damage it has done to the news industry, enables the knowledge-claims of legacy news companies to be shared quickly, broadly, and freely to a degree that could never have been matched by the technologies

website's bearing a resemblance to something produced by a professional news organization, rather than heuristics more appropriate to the digital age (Steinmetz 2018).
179See, for instance, Goldman (2009)

that preceded it. Since the internet and social media both facilitate the spread of massive amounts of scientific and professionally newsworthy information, it is perplexing when commentators like Schiffrin claim that the internet *as a whole* does not appear to the standards of professional journalism and science and researchers at Pew and YouGov-Cambridge – both of which organizations have emphasized the special democratic importance of professional journalism – frame their questions in a way that insinuates that the trustworthiness of social media content ought to be summed up in sweeping generalizations.

As with all the other criteria discussed so far, when commentators neglect to demonstrate an appreciation for and attempt to take into account the considerable *causal complexity* entailed in their discussions of the epistemic well-being of democracy, they don't just risk making mistakes in their diagnosis of that well-being, but risk reaching a diagnosis that may be the *exact opposite* of what is truly the case. For instance, while it has certainly undermined the advertising revenue model that subsidized the legacy news industry for over a century and may have contributed to declines in newsroom staffing, particularly for local periodicals, it is still not clear that social media has decreased the *epistemic influence* of professional news organizations *as a whole* even so. Just because such papers' ad revenues and number of paying subscribers have gone down does not mean that there has been a decline in the number of people who *gain* and *make use of* the knowledge disseminated by these entities, or what might be called the effective audience of those organizations. The *effective audience* for professional news is not currently, and has never been, limited to its regular viewers or subscribers. Instead, the public's awareness of news has always depended in large part by its being re-transmitted through interpersonal relationships between lay citizens. As Elihu Katz wrote in 1957, the simplified “image of the audience of a mass of disconnected individuals hooked up to the media but not to each other” is not an apt description of the relationship between the news and democracy.¹⁸⁰ What Katz was pointing out was that in 1957 one did not need to be a subscriber to *The New York Times* or *Popular Science* to gain and make use of the knowledge disseminated through those publications. Having chatty friends who read them regularly, or even being several degrees removed from a reader of such publications, may well have sufficed.

¹⁸⁰Katz (1957), 61

Katz was describing how our embeddedness in *social networks* can greatly amplify our ability to gain and make use of the knowledge, even when we do not gain that knowledge first-hand. I will call knowledge gained through one's social network in this way *second-hand knowledge*. While second-hand knowledge has always been integral to most individuals' understanding of current events and science (as well as many other topics),¹⁸¹ the internet, mobile technologies, and especially social media have *massively* expanded our ability to make use of social connections in ways that allow for the effective transmission of knowledge of this sort. Nothing is easier than linking a friend to an article or sharing a summary of its gist through text, posting a link containing its headline and summary text on one's Facebook wall, or Tweeting or re-tweeting an article or a description of it by some professional journalist him/her self on one's Twitter account. While this fact has often been bemoaned by those who see it as conducive to the spread of misinformation, it enables the spreading of *true* information to the same degree. Thus, it is potentially easier than ever for *non-subscribers* and *irregular news viewers* to nevertheless gain and make use of the knowledge produced by professional news organizations. In fact, so much easier is it to re-transmit information in the aptly named “information age” that it is possible that the overall epistemic impact of professional news organizations within society is *increasing* even if the absolute number or even percentage of the total public that *subscribes to* or *personally reads/visits* the information disseminated by these sources declines, and is doing so precisely *because of* social media. This is important because, as I have argued, the evident concern of commentators who mourn the decline of the legacy news industry stems from their belief that the knowledge disseminated by that industry contributes to the epistemic well-being of democracy. This concern for *democracy* is distinct from other concerns that might more legitimately cause one to mourn the rise of social media, such as the loss of career prospects for individual reporters or losses in advertising profits for legacy news companies. If our concern is about the epistemic well-being of democracy and our hope is that its citizens' beliefs should be informed by the professional journalism industry, then it is not the *immediate* but rather the *effective* audience that we ought primarily to be worried about.¹⁸² Indeed, it is very likely that the

¹⁸¹In fact, every contemporary scholar I am aware of who has commented on the matter argues that our reliance on what I have here called “second-hand knowledge” extends to very nearly every one of our beliefs. See Hardwig (1985) and Coady (1992) for just two examples.

¹⁸²It is not, however, the *only* thing we might be worried about. We might also be worried about the ways in which competition from “click-baiters” pressures professional news organizations to engage in similar tactics,

effective audience of well-respected news organizations like the *Times* and the *Washington Post* have *grown* in the digital age, especially when one compares digital-age developments against the most plausible counterfactual scenarios. Now, it is true that the internet has had other effects on the legacy news media, such as dramatically undercutting its advertising revenue, and has contributed to alterations in news-gathering practices and an overall reduction in the number of journalists, especially local journalists. Surely some of these developments have reduced the financial viability of news sources and undermined their ability to provide high-quality news. My point is not that the internet and social media have clearly improved and aided legacy journalism but that the relationship between 21st century communications technologies and legacy journalism is complex enough to make an overall assessment on that industry difficult.

Adding to the complexity of assessing 21st century communications technologies' impact on the legacy news media and on democracy in general is that *different aspects* of those technologies are likely to produce *different effects*. The likely epistemic impact of the social media accounts of scientists like Neil DeGrasse Tyson and legacy news organizations like the *New York Times*, for example, are likely to be far more beneficial than that of social media accounts managed by the Flat Earth Society. Another source of complexity is that the *same aspect* of 21st century communications technologies may have *multiple, different* effects, as is the case for the “share” button on Facebook, which allows both articles from legacy news organizations like *The Washington Post* and fake news articles from the (now-defunct) *WTOE 5* website to spread with unprecedented ease and speed. This complexity extends to the different ways different users of these technologies process the information they find on them as well as the different ways the *same* user uses the *same* technologies at different times and states of mind. When Progressives over-generalize about the average quality of the information available on “the internet” and “social media,” they fail to take this complexity into account. As I have explained in this sub-section, one consequence of their failure to do so is that they tend to overlook the ways those technologies may be beneficial to the epistemic well-being of

decreasing the quality of the information published by professional news organizations themselves (see Davies 2008 and Silverman 2015). But given the perennial complaints of news critics about “sensationalism,” it is not clear that newspapers have *ever* been very immune from pressures to oversimplify and sensationalize in order to gain viewers' attention. Beyond this, Davies (2008) puts together a compelling case that most of the economic pressures that may be argued to have decreased the epistemic quality of the contents of newspapers began in the 1980s and 1990s, long before either the internet or social media had much of an impact on the industry.

democracy *even with regard to* the very institutions they view as being challenged by those technologies, such as the “community of science” and the legacy news media.

But another consequence, which I find more troubling, is that Progressives' tendency to generalize about 21st century technologies in sweepingly negative terms encourages citizens to take an approach to those technologies that is unlikely to allow them to make the most of their potential and may be damaging to democracy both epistemically and socially. Epistemically, Progressives' habit of characterizing “the internet” and “social media” as bastions of “unreliable” information encourages citizens to approach the information they encounter via those technologies with an air of suspicion while approaching the contents of scientific journals and legacy news broadcasts (for example) with a disposition of credulity. As recent polling data such as the Pew study cited earlier demonstrate, the public is rapidly adopting this attitude of suspicion toward social media in particular, with over half reporting that they *expect* the news they find on social media *to be false*. The very phrasing of this query encourages citizens to lump all the contents they find on social media into a single category and attach a label of (un)reliability to it. While many see the encouragement of this kind of suspicion toward the contents of social media as the adoption of a healthy attitude of skepticism toward a sketchy medium, I disagree. For the many citizens whose primary or only mode of access to the claims of scientists and the legacy news media *is* the internet, adopting this attitude of suspicion would deter them from adopting the very beliefs Progressives believe they ought to adopt in order to foster the epistemic well-being of democracy. But more worrisome than this danger are the epistemic and social implications of this attitude of suspicion for lay citizens' views of *one another*. Much of the content citizens encounter on social media is posted by other citizens, so to imply that citizens ought to view the contents of social media with suspicion is to imply that they ought to view their fellow citizens with suspicion, that is, to view them as either dishonest or lazy or possessing poor judgment, and accordingly to view them as if they cannot be depended on to pass on reliable information. In other words, it encourages citizens to view one another in the same pessimistic light in which lay citizens are viewed by Progressives. I have suggested in this chapter, and will go on to argue more forcefully throughout this dissertation, that this pessimism is unjustified. To the degree that it *is* unjustified and lay citizens *can* be trusted to make reliable judgments and pass on reliable information online, encouraging them to view each

other's claims with suspicion stifles the spread of truth. Additionally, to the degree that citizens adopt this suspicion and skepticism toward each other while simultaneously adopting the disposition of credulity toward established epistemic institutions in the way encouraged by Progressives, they become increasingly dependent for their beliefs on the reports disseminated by those institutions and decreasingly receptive to challenges to those reports that disseminate from those fellow citizens. This undermines the primary means citizens have of discovering and resisting illegitimate authority of any kind, epistemic or otherwise. Beyond these epistemic concerns, political scientists and social theorists in recent decades have convincingly argued that for a variety of reasons “generalized trust” between citizens is a vital element of a well-functioning democracy.¹⁸³ Encouraging citizens to view one another as some combination of dishonest, lazy, incompetent, and otherwise untrustworthy purveyors of information is very likely to discourage the development of this sort of generalized trust, a problem all the more worrisome in a time such as the present which is characterized by high levels of affective polarization. For these reasons, I believe the attitude of suspicion toward 21st century communications technologies encouraged by Progressives' tendency to sweepingly generalize their contents as “unreliable” poses both an epistemic and a social threat to democracy. This gives us good reason to oppose the Progressive tendency to generalize about the contents of “the internet” and “social media” in sweepingly negative terms, at the very least until that generalization is based on a more intricate and nuanced model of the likely causal relationship between those technologies and the epistemic well-being of democracy.

I will close this sub-section by suggesting that, just as was the case in the previous sub-sections, Progressives' pessimistic interpretation of the epistemic impact of 21st century communications on technologies depends heavily on their having *presupposed* a deeply pessimistic view of the judgmental capacities of lay citizens. It is easy to see that the internet and social media host a wide variety of contents and that those contents include the claims of members of the “community of science” as well as the legacy news media. It should also be obvious that different people use those technologies in different ways and in different states of mind. Given this state of affairs, it is not immediately obvious why we should expect the “average” quality of information on social media should be low or the overall epistemic impact

183 E.g. Warren (1999), Putnam (2000), Dekker and Uslaner (2001)

of 21st century communications technologies on democracy to be negative. These mysteries are readily resolved, however, and Progressives' habit of generalizing in sweepingly negative terms about the epistemic impact of 21st century technologies on democracy made sense of, if we simply presume that lay citizens are some combination of dishonest, lazy and incompetent most of the time. If this is the case, the presence of members of the “community of science” and the legacy news media online will not matter much epistemically, since they may be expected to be largely ignored or disbelieved by a beleaguered public. If the judgmental capacities of the public can be adequately characterized in this way, there is less need to develop a complex causal model linking social media to the epistemic well-being of democracy, since the assumption of largely unreliable judgment on the part of the citizen means that the variety of the inputs (the types of informational sources and modes through which citizens encounter those sources online) does not matter much (since citizens will make poor judgments about what to believe) and there is no meaningful variety in terms of the informational system's outputs (citizens' beliefs and “shares”), as they are assumed by the model to be largely misguided. Thus, once again, Progressives' neglect of one of the criteria I have identified as necessary to properly assess the epistemic well-being of democracy can be elegantly explained by their having entered into the discussion having already *presupposed* a deeply pessimistic view of the judgmental capacities of lay citizens – a presupposition that is not clearly justified by either present evidence or historical trends. This pessimistic starting point leads Progressives to pessimistic conclusions about the epistemic impact 21st century technologies are having on democracy, which in turn reinforce their pessimistic view of citizens.

Conclusion: Tension Between the Progressive View and the Ideals of Liberal Democracy

In the Introduction and Chapter One of this dissertation, I argued that deep pessimism about the judgment of lay citizens is incompatible with any compelling defense of liberal democracy and signaled my intent to argue both that this deep pessimism is unjustified. In this chapter, I have shown a few specific ways I believe Progressives' pessimistic conclusions about the epistemic impact of 21st century communications technologies on democracy are at best premature and at worst fundamentally misguided. In each case, those pessimistic conclusions

were not well supported either by present evidence or by historical trends, and in each case Progressives' propensity to rush to those pessimistic conclusion could be made sense of by assuming a deeply pessimistic view of the judgmental capacities of lay citizens.

In addition to their deep pessimism about the judgmental capacities of lay citizens and the epistemic impact of 21st century communications technologies on democracy, one more commonality in all the critiques I have leveled against Progressives here is worth pointing out. This is the dominant role played by *established epistemic authorities* in Progressives' implicit view of the proper epistemic functioning of democracy and the simultaneously subordinate role played by lay citizens. The Progressive View emphasizes the democratic importance of citizens *trusting* the claims of the “scientific community” and the legacy news media, but Progressives rarely emphasize the need for lay citizens to *question, criticize* or *resist* those claims. This makes perfect sense given Progressives' deep pessimism about the judgmental capacities of lay citizens. After all, if lay citizens are the kind of exceedingly poor judges of truth Progressives often depict them as being, there would seem to be little value – and possibly much harm – in encouraging them to question, criticize and resist the claims of experts and professional knowledge-seekers on the basis of their own judgment. As I argued at multiple points in this chapter, Progressives' conclusions about what best serves the epistemic well-being of society as a whole make sense if one accepts their deeply pessimistic starting point with regard to lay citizens' judgment. What makes *less* sense given this deeply pessimistic view of the judgmental capacities of lay citizens is Progressives' avowed support for *liberal democracy* as defined in Chapter One. If the main thing that determines the epistemic well-being of our society as a whole is citizens' willingness to place their trust in certain established epistemic authorities like the “community of science” and the legacy news media and citizens cannot be trusted to reliably assess the trustworthiness of those authorities using their own powers of judgment, why would Progressives continue to support protections for free speech and universal, equal suffrage? The combined effect of these two policies would seem to encourage precisely those conditions Progressives most fear – the ability of lay citizens to access and disseminate a wide variety of information from all sorts of sources of varying quality and the encouragement of those citizens to *decide for themselves* which information to believe and which leaders to elect and which policies to support on the basis of that information.

Progressives who view lay citizens' willingness to trust certain established epistemic authorities as integral to the epistemic well-being of democracy are forced on the horns of a dilemma. Either they must abandon their support for liberal democracy or they must amend their pessimism about the judgmental capacities of lay citizens. If lay citizens trusting in the proper epistemic authorities is as integral to the epistemic well-being of a society as Progressives suggest, it is not something that should be left in the unreliable hands of a public whose judgment is critically compromised. Theorists who care as much about the epistemic well-being of a society as Progressives seem to should therefore only embrace liberal democracy if they view lay citizens as possessing good enough judgment that they can be largely trusted to come to believe the truth and trust the appropriate epistemic authorities over time. If they think that judgment is exceedingly poor, they should abandon either their commitment to *liberalism* or their commitment to *democracy*. As we shall see later, the former is precisely what Progressives have increasingly begun to do. The argument of this chapter suggests this move is at best premature and at worst deeply misguided. Progressives' deep pessimism about the judgmental capacities of lay citizens is not clearly justified, nor do either current evidence or historical trends support their pessimistic conclusions about the epistemic impact of 21st century communications technologies on democracy. On the contrary, the evidence and arguments considered in this chapter suggest we have every bit as much reason to be *optimistic* about the epistemic impact of 21st century communications technologies on democracy as we do to embrace the deep pessimism favored by Progressives. Accordingly, at least so far, we have every bit as much reason to resolve the Progressive dilemma in the opposite way from that favored by contemporary Progressives – that is, by viewing the judgmental capacities of lay citizens in a more optimistic light and treating citizens as dependable enough judges of truth to justify the protections of free speech and universal, equal suffrage that are the key elements of any liberal democracy worth defending. By the end of this dissertation, I hope to have made a compelling argument that this is precisely what we should do, but for now I just hope to have begun to drive the wedge of suspicion into readers' minds about the currently rampant negativity in our interpretations of the epistemic impact of 21st century communications technologies on democracy and the deeply pessimistic view of the judgmental capacities of lay citizens it implies.

The argument of this chapter is meant to be an *immanent critique* of the Progressive View. That is, my critique is meant to undermine Progressive pessimism about lay citizens and the epistemic direction of democracy *even if* one accepts the Progressive model of democratic epistemic well-being. This is not because I agree with the way Progressives conceptualize the epistemic well-being of democracy – far from it. In fact, the rest of the dissertation will be dedicated to criticizing many aspects of this model – aspects I believe not to be unique to Progressives but, to the contrary, to be characteristic of prevailing views of what it means for democracy to be “doing well” epistemically. In each case, I will connect the epistemic behaviors of *individual citizens* to their likely contribution to *collective epistemic outcomes*. I will argue that in many ways, the epistemic behaviors encouraged by Progressives are likely to be detrimental to the epistemic well-being of democracy while the behaviors they consistently deride are likely to contribute to that epistemic well-being, and that the same is true of the public policies Progressives support on the grounds of protecting the epistemic well-being of democracy. One reason this is so is because individual behaviors can lead to counter-intuitive collective results. That is, what may plausibly be viewed as *epistemic vices* at the individual level can, when practiced by many individual members of a collective, turn out to be *epistemic virtues* for a society. This will be the argument of Chapter Three.

Chapter Three: Individual Vice, Collective Virtue

The preceding chapter introduced the notion of the epistemic well-being of democracy. In it, I identified a number of problems with the ways Progressives tend to evaluate democracy's epistemic well-being in the present. Because of these problems, I argued, Progressives' negative evaluations of 21st century communications technologies' impact on the epistemic well-being of democracy are at best premature and at worst fundamentally misguided. Throughout the chapter, I emphasized the role played by Progressives' having *presupposed* a deeply pessimistic view of lay citizens' judgment. In each of the examples mentioned, Progressives' propensity to rush to pessimistic conclusions about the direction of democracy and 21st century technologies' influence on it could be neatly explained by their having adopted a view of lay citizens as possessing deeply unreliable judgment. I closed the chapter by pointing out this deep pessimism's *affinity with* Progressives' heavy emphasis on the need for lay citizens to place their trust in epistemic authorities and its *dissonance with* support for the sort of liberal democracy defined in Chapter One. Since neither contemporary empirical evidence nor historical trends support the level of pessimism about lay citizens' judgment that currently reigns among Progressives, I argued, we need not feel compelled to join them in either respect. So far, at least, we have seen no compelling reason either to abandon support for liberal democracy or to follow Progressives in so heavily emphasizing the need for lay citizens to defer to established epistemic authorities, such as the legacy news media and the “community of science,” as a necessary condition for preserving the epistemic well-being of democracy.

Recall that in Section One of Chapter Two I identified what I called the *Progressive View* and the *Progressive Model of Democratic Epistemic Well-Being* built upon that view. In a nutshell, I argued that Progressives view democracy's epistemic well-being as built primarily on (1) citizens having access to and believing (i.e. acting on the basis of) *reliable information*, where (2) *reliable information* is defined as information disseminated by the relevant professional or expert source(s). While I mentioned some reservations about this as a set of

normative standards, I did not dwell on them at any length. Instead, what I wrote in Chapter Two was chiefly intended as an *immanent* critique of Progressivism, one that accepted the Progressive View the model of democratic epistemic well-being built on it. In the next few chapters, my critiques of Progressivism will expand beyond these confines. In this chapter, I will argue that even if citizens are afflicted with the sorts of *epistemic vices* Progressives often impute to them, in many cases those very vices actually make democracy *as a whole* better off epistemically. In the next two chapters, I will move on to critique Progressives' deferential view of responsible democratic citizenship. My goal is to argue that the Progressive perspective that so dominates scholarly and mainstream assessments of contemporary democracy and the epistemic impact of 21st century communications technologies on it is both incompatible with liberal democracy and deeply misguided, and that therefore instead of endorsing the Progressive View, we ought to look elsewhere in seeking to define what it means for democracy to be doing well epistemically and of the sorts of civic dispositions and practices likely to contribute to that well-being.

This chapter will consist in five sections. In Section One, I will set up a framework of individual and collective epistemic vices and virtues by relating those concepts both to contemporary conversations about the epistemic impact of 21st century communications technologies as well as certain instances in the history of science. The main point I wish to establish in this section is that many behaviors that can sensibly be dubbed epistemic vices at the individual level ought to be considered virtues at the level of the collective. The next three sections will use this framework to raise objections to Progressives' view that 21st century communications technologies are contributing to the proliferation of collective epistemic vice. In each section, I will show how the behaviors of individuals condemned as threats to the epistemic well-being of democracy by Progressives may plausibly be contributing to its epistemic well-being. In Section Two, I will argue that motivated reasoning and other cognitive biases, whose proliferation is often cited as evidence of citizens' unfitness for their democratic role, often actually benefit our collective acquisition of knowledge. In Section Three, I will defend anti-vaccers and similar communities who dissent from the scientific mainstream as epistemic assets to, and signs of the robustness of, the epistemic well-being of contemporary democracy, as they are valuable sources of vigilance and critique toward authorities seldom questioned by the rest of us. In Section Four, I will argue that the “echo chambers” and “filter bubbles” cited by

Progressives as signs of technology-fostered ignorance and laboratories for mutual antipathy also serve as ways for “deviant” communities to find mutual support and develop alternative lifestyles and approaches to politics that can be of great value to liberal democracies. In Section Five, I will argue that recognition of the potential benefits of these phenomena deemed suggests that we have more to gain than lose by encouraging citizens to assertively exercise doxastic self-rule, even though Progressives are right that too much assertiveness on the part of citizens can be vicious and lead to some bad outcomes. Because the benefits outweigh the risks, I will conclude that we ought to encourage citizens to err on the side of assertiveness when it comes to exercising doxastic self-rule rather than err on the side of timidity and deference, as Progressives encourage them to do, and that our approach to public policy ought to be such that we give citizens the chance to make effective use of this assertiveness. In practical terms, this means opposing the sorts of censorship and top-down management of citizens' informational environment supported by Progressives.

Section One: Individual vs. Collective Epistemic Vices and Virtues

Democracy, Liberty, and Epistemic Vice

Progressives clearly view lay citizens of 21st century democracies as riddled with epistemic vices. Although some of the language contemporary scholars use to describe these vices, such as the terms “cognitive bias” and “motivated reasoning,” is unique to our age, this pessimistic view about lay citizens is nothing new. For well over two millennia, lay citizenship in democracy was associated with epistemic vice, and for this reason (among others) democracy was viewed as an inferior and unstable political system. In the *Republic*, “Socrates”¹⁸⁴ famously depicts the democratic citizen as lawless, morally relativistic and hedonistic.¹⁸⁵ For Plato, as for others, the aspect of democracy that facilitated vice and allowed it to wreak its havoc was freedom. In the eyes of critics like Plato, the free and open environment so treasured by democratic citizens would only make society vulnerable to the misuse to which it would

¹⁸⁴ As Karl Popper has forcefully argued, the “Socrates” who serves as Plato's mouthpiece in the *Republic* espouses views so diametrically opposed to those most closely associated with the historical figure of Socrates that the two ought to be deliberately distinguished, hence the quotes around the name “Socrates” who appears as a character in this and several of Plato's other dialogues (Popper 1971, 131-2).

¹⁸⁵ *Republic* (ch. 8)

inevitably be put by the many, creating an atmosphere unfriendly to truth, order, and consequently good governance. This view, connecting democratic citizenship with freedom, freedom with epistemic vice, and vice with social instability and undesirable political outcomes, fit well with the forms of politics – feudal aristocracy, hereditary monarchy and the hierarchical Catholic church – that dominated Europe for the next two thousand years, and was still the dominant view by the time of the American Revolution.¹⁸⁶

Though practically no one, including Progressives, would publicly claim not to support what is colloquially called "democracy" today – that is, roughly speaking, a political system in which political leadership and certain fundamental political decisions are decided upon by popular elections characterized by universal suffrage – for many, the association of democratic citizenship with vice via the vector of freedom remains. Like Plato and Montesquieu, many seem to think the sort of epistemic virtue required to make democracy succeed is largely lacking in the citizenry; that whatever virtues the few may possess are overwhelmed by the vice of the many; and that these vices are caused and/or exacerbated by citizens' access to an excessive array of freedoms. This is the underlying theme of much of the work examining the impact of 21st century communications technologies by contemporary political science and communications scholars. Take, for example, the work of prominent scholars Cass Sunstein and Markus Prior. Both have championed arguments that implicate the increased freedom of choice afforded to citizens by the decreased cost of and increased access to information technologies as causes of ideological segregation, polarization, and other of what Prior and Sunstein, along with plenty of their contemporaries, view as democratic ills. For these thinkers and their numerous sympathetic colleagues, the route through which 21st century communications technologies mete out their pernicious effects on democracy runs straight through the vices of citizens, whose shortsightedness and assorted "cognitive biases" allow this expansion in freedoms to produce in our society ideological segregation, polarization, and other of what Prior and Sunstein, along with plenty of their contemporaries, view as democratic ills.¹⁸⁷

186 For instance, in the mid-18th century Montesquieu wrote that the sorts of virtue required to make democracy work, including epistemic virtue, was highly unnatural and unlikely to be attained by citizens under conditions of freedom. In an effort to build up that virtue, Montesquieu called for democracies to institute intensive and obligatory education, censorship, enforced frugality, and highly restricted property rights (Montesquieu 1952 [1748], book 5 chs. 4-5).

187 Prior (2007), Sunstein (2016)

Individual vs. Collective Vices and Virtues

I believe there is a gap in the logic that frequently underlies such arguments. That gap appears when thinkers attempt to scale up their observations of individual-level behaviors to collective outcomes. My central argument is that an increase in epistemic vice among *individuals* does not necessarily signify a decrease in the overall epistemic performance of the *collective*. Even if we grant that democratic citizens are, on average, as riddled by these vices as commentators suppose, and that increases in ideological segregation, polarization and "extremism" so commonly lamented by Progressives are in fact undesirable, we need not conclude from their mere presence vicious consequences for *democracy* as a whole. In fact, I will argue that what may be fairly dubbed epistemic "vices" at the individual level are often *virtues* when considered at the level of the collective, and that this is likely to be the case for several of the phenomena most often lamented by pessimists about 21st century communications technologies.

Since I am going to be using the term frequently, it seems only polite for me to begin by saying what I mean by the term "epistemic vice." I take "vice" to mean a bad habit, where "bad" can mean either immoral or simply destructive to one's purposes. And I add the adjective "epistemic" to signify that the vices I'm talking about have to do with how citizens go about acquiring and/or spreading knowledge. An "epistemic vice," then, is a habit related to the acquisition and/or dissemination of knowledge that is either morally wrong or destructive to the purpose of acquiring and/or spreading knowledge.

21st Century Communications Technologies as Vectors for Epistemic Vice

Though the term is seldom used outside of philosophy departments, the presence of epistemic vice among the lay citizenry is, in one way or another, cited as a grounds for worry about the epistemic well-being of democracy in a wide range of current discussions in the academy and the mainstream. Let me offer three examples, each taken from a different discursive arena, to give some sense of the topic's breadth.

Example one can be found in the pages of journals in the fields of psychology, economics, political science, and others, where the notion of "cognitive biases," prominent among which is the concept of "motivated reasoning," has become ubiquitous. The gist of much

of this literature is that citizens' thinking is riddled with biases that lead them to adopt beliefs that are systematically distorted.¹⁸⁸ Though some scholars in this tradition have sought to use their knowledge of the normal workings of cognitive biases and basic human psychology to develop theories about how to overcome or at least mitigate them,¹⁸⁹ in many cases the prevalence of these biases is cited as a grounds for pessimism about democracy. Robert Shapiro and Yaeli Bloch-Elkon sum up this line of reasoning nicely. After making the case that partisan motivated reasoning inhibits citizens' ability and/or willingness to believe the truth, Shapiro and Bloch-Elton conclude that when "partisan disagreements about important factual issues show that large subsets of the public are necessarily wrong about the facts, then there is clearly cause for concern."¹⁹⁰ The biases cited by Shapiro and Bloch-Elkon and others can be thought of as epistemic vices, and the literature that treats them as grounds for pessimistic conclusions about the state or current direction of democracy as arguing that to the extent that these vices prevail among citizens, democracy is bound to suffer, at least epistemically.

A second contemporary conversation that treats epistemic vice among the public as potentially problematic for democracy revolves around what Tom Nichols calls the "death of expertise,"¹⁹¹ i.e. certain groups of citizens' insistence on rejecting the conclusions of experts and/or professionals in favor of their own. Widely discussed examples of this behavior revolve around subsets of citizens' opposition to the scientific consensus on anthropogenic climate change and the necessity and effectiveness of vaccines, as well as English citizens' apparent indifference to economic experts' views on the likely effects of England's exit from the European Union, or "Brexit." In each of these cases, groups of lay citizens have been heavily criticized by numerous commentators for what I will call *epistemic hubris*, the choice to rely upon their own judgment in an arena in which they possess little to no expertise over that of those who do possess it, and increases in such hubris among the citizenry are treated as straightforwardly detrimental to the epistemic well-being of democracy.

A third example of a contemporary line of conversation in which the epistemic vices of citizens are often lamented as a detriment to democracy was referred to briefly in the

188 E.g. Taber and Lodge (2013), Brennan (2016)

189 Lewandowsky et al (2012)

190 Shapiro and Bloch-Elkon (2008), 131

191 Nichols (2017)

introduction. This is the conversation that has arisen over the uses to which 21st century communications technologies, such as the internet and social media, have been put by lay citizens, and the impact of that usage on democracy. Two of the most prominent voices in that conversation are those of Markus Prior and Cass Sunstein. Prior connects the increased choice in information and entertainment sources made available by technological advances in telecommunications to political polarization and participatory inequality.¹⁹² Meanwhile, Sunstein joins Prior in his worries about the aggregate effects of a high-choice media environment and warns that its concurrence with the rise of algorithmic sorting and the decline of "general-interest intermediaries" like the professionalized news of the broadcast era are likely to lead to the development of "echo chambers" that isolate groups holding different worldviews from one another and thereby foster ignorance and extremism.¹⁹³ For both thinkers, natural and deep-seated shortcomings of the human mind, such as a preference for the company and claims of those who share one's own beliefs – in other words, epistemic vices – play a pivotal role in producing the political outcomes they lament. Like the psychological literature on "cognitive biases," this line of thinking has become highly influential both across academic disciplines and in the mainstream, with pundits and everyday citizens frequently lamenting the decline of the professional news media and accusing their "echo chamber"-confined compatriots of the same combination of laziness, short-sightedness and moral failure implicit in the writings of Prior and Sunstein. And, once again, the straightforward implication is that to the extent such behaviors prevail among citizens, the epistemic well-being of democracy as a whole will suffer.

Implications for Liberal Democratic Citizenship

In these contemporary conversations and others related to them, one can spot a common chain of reasoning that connects epistemic vice on the individual level to harms to democracy on the collective level. It is the prevalence and often unqualified acceptance of this chain of reasoning I wish to critique in the remainder of this essay. For it is my belief that even if we accept these commentators' characterization of the epistemic behaviors of democratic citizens as "vices" at the individual level, no straightforward inference can be made about the effects on

¹⁹² Prior 2008

¹⁹³ Sunstein 2016

democracy of millions of individual citizens or thousands of small groups practicing these very same behaviors. On the contrary, I hope to show that many of what we commonly, and perhaps correctly, call epistemic vices when practiced by individuals are often virtues at the level of the body politic, and to argue that this *is* in fact the case for at least some of the most commonly lamented vices of the present.

The argument of this chapter has important implications not just for the ways we assess the epistemic well-being of democracy and 21st century communications technologies' impact on it, but also on the public policies we endorse and the dispositions and behaviors we encourage in citizens in the name of fostering that epistemic well-being. The Progressive line of thinking subscribed to by most scholars and mainstream commentators in the conversations mentioned above (a) casts strong doubt on whether citizens are fit for liberal democracy as defined in Chapter One, (b) encourages a deferential attitude toward experts while discouraging lay citizens from – indeed, often disparaging them for – exercising doxastic self-rule, and (c) generates opposition to the free flow of information and opinion between lay citizens online, since an appreciable amount of the time such an exchange is likely to lead judgment-deficient citizens to endorse disinformation and wrongly reject and devalue the claims of experts. I believe all of these are mistakes that stem largely from deficiencies in the ways Progressives tend to think about what it means for democracy to be doing well epistemically and the sorts of dispositions and behaviors of lay citizens that are likely to contribute to that well-being in the long haul. One of those deficiencies is the tendency for Progressives to assume that behaviors that are detrimental to an *individual's* likelihood of forming reliable beliefs, i.e. *individual epistemic vices*, also detract from the epistemic well-being of the *collective*, and therefore ought to count as *collective epistemic vices*. Rejecting that position and exploring the implications of that rejection for citizenship and public policy are the main tasks of the rest of this chapter.

Moralistic vs. Descriptive Vices

By the end of this chapter I will have argued that behaviors properly viewed as vices at the individual level may prove to be virtues when viewed at the level of the collective. Stated that way, this view is nothing new. It is similar to the concept of *Private Vices, Publick Benefits* that served as the alternate title for Bernard Mandeville's *Fable of the Bees*, which scandalized

early 18th century Europe by suggesting that private vices created social benefits.¹⁹⁴ An analogous line of argument can be found in the famous *Wealth of Nations* passage in which Adam Smith claims that it is "not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own self-interest."¹⁹⁵ The constitutional thought of American founding father James Madison, who sought to design a system of government that used two attributes commonly viewed as vices in republican thought – selfishness and personal ambition – to keep power-seekers beholden to the public interest, seems likewise to have been influenced by this vein of thought. In the thinking of both Smith and Madison, a quality that may be fairly described as a vice when practiced by an individual at the local level is, once integrated into the machinery of society, viewed not only as having *some kind* of collective benefit but indeed to be a manifestation of its *exact opposite* virtue. That is, rather than resulting in outcomes resulting in the individual's gain at the public's expense, the individual vices mentioned by Smith and Madison are construed as ultimately resulting in the public's *benefit*.

These venerated examples are familiar and are often taken into account as by political theorists. However, as illustrated above, the conception of "vice" in these analyses is typically *moralistic*, while the collective outcome of concern contains a *descriptive* element. That is, for Smith and Madison, "vice" means simply a failure to live up to some moral code, while the articulation of a collective benefit said to result from that vice entails a *descriptive* claim. The genius of Smith and Madison was to draw out the implications of this line of thinking for public policy, but even by their time the idea that a morally vicious act might in fact result in desirable outcomes was old hat, with Machiavelli having published *The Prince* more than two centuries before. While these themes in Machiavelli, Smith and Madison are very familiar to political theorists, much less consideration has been given to the possibility that a similar divergence might apply to the relationship between collective-level outcomes and the individual-level practice of what I will call *descriptive* vices, i.e. actions or habits that are detrimental to the goal(s) of the agent, be they morally praiseworthy or no. At the very least, this possibility has not been explored to any appreciable depth by commentators interested in the epistemology of democracy. Instead, the most common way in which political thinkers connect individual-level

¹⁹⁴Mandeville (1962 [1714])

¹⁹⁵ Smith (1937 [1776]), Bk. 1, Ch. 2

epistemic behaviors to collective-level outcomes is by recommending some set of educational arrangements for maximizing the epistemic performance of individual citizens¹⁹⁶ or the adoption by citizens at large of some epistemic disposition thought to be beneficial to knowledge-seeking when practiced by individuals.¹⁹⁷ In both cases, the relationship between individual and collective epistemic vice/virtue is presumed to be uniformly positive, that is, greater the individual epistemic virtue, the better off the epistemic well-being of the collective is presumed to be. The assumption seems to be that because better education and/or a more learning-conducive knowledge-seeking disposition may plausibly be expected to improve *each individual's* ability to seek/acknowledge the truth, we should expect better education and/or the adoption of a more learning-conducive knowledge-seeking disposition by more citizens to improve the epistemic condition of *democracy* as a whole.

Using two essays from contemporary philosophy of science to illustrate my logic, I will contest this assumption. Like Smith and Madison, I will argue that the vicious habits and actions of individuals can produce collective benefits. However, unlike Smith and Madison I will explicitly insist that this principle extends beyond the realm of individual *moral* vice and into the realm of what I have called individual *descriptive* vice in the realm of epistemology. That is, even epistemic actions and/or habits that are detrimental to individual knowledge-seeking can in principle, and in some cases do in reality, serve to improve collective epistemic well-being.

Example One: Phlogiston and Plate Tectonics in Philip Kitcher's "The Division of Cognitive Labor"

The first essay I will discuss is Philip Kitcher's "The Division of Cognitive Labor." Kitcher starts off this essay with the historical example of phlogiston theory, which, Kitcher reports, despite having been "favored by almost every chemist in Europe" in the 1770s, was

196 Democratically disposed thinkers' preoccupation with education was pithily summed up by Walter Lippmann nearly a century ago when he wrote that "education has furnished the thesis of the last chapter of every optimistic book on democracy written for one hundred and fifty years" (Lippmann 1925, 22). While fewer political scientists may be as optimistic as the democratic theorists of this era, those who emphasize education's importance to democracy usually do so on the tacit assumption that maximizing educational attainment among citizens is bound to positively correlate with desirable political outcomes.

197 The most famous proponent of the theory that democracy depended on the adoption of a suitable epistemic disposition among the citizenry was, of course, John Dewey, whose arguments in this vein have recently been championed by Elizabeth Anderson (Anderson 2006).

"dead" by 1804, mostly due to the work of Antoine Lavoisier. While over two centuries of hindsight indicate that abandoning phlogiston theory was a good move for chemists to make, by the point in time referred to by Kitcher both phlogiston theory and Lavoisier's "new chemistry" were backed by numerous experimental results, conducted over a long span of time and backed by numerous theories whose results seemed persuasive – albeit in favor of different conclusions – to large numbers of inquirers. Even granting that the experiments reported by Lavoisier between the 1770s and 1800s tilted the evidentiary balance away from phlogiston in favor of Lavoisier, Kitcher asks whether the "initially uniform opinion, sudden jumping of ship, and [rapid settlement upon a] new consensus" over the span of a scant 30 years should really be interpreted as a positive indicator of the epistemic condition of the field of chemistry as a whole.¹⁹⁸ Kitcher argues in the negative, suggesting that a better indicator of the epistemic condition of the community of chemists may have been Joseph Priestley, who continued to defend phlogiston theory even after it had been abandoned by nearly every chemist in Europe. Importantly, Kitcher's criticism of the scientific community and favorability toward Priestley in this instance has nothing to do with the idea that phlogiston theory was as *plausible* as Lavoisier's "new chemistry," even from the perspective of the scientists at the time. On the contrary, Kitcher grants that at a certain point the weight of the evidence against phlogiston theory probably made it *individually* rational for any well-read European chemist, including Priestley, to view it with skepticism. Nevertheless, Kitcher argues that even though it was "unreasonable" for Priestley "to persist [defending phlogiston theory] as long as he did...from the point of view of the *community of chemists*, it was no bad thing that Priestley (and a few others) gave the phlogiston theory every last chance."¹⁹⁹

Why does Kitcher argue that it is desirable from a scientific perspective, at least in some instances, for some subset of individual scientists to irrationally resist the conclusion best supported by the evidence? The answer can be found in the the next historical example he cites, that of Alfred Wegener's theory of continental drift. In the 1920s and 1930s, Wegener's theory that the continents were once joined together but eventually broke up and drifted to their current positions "seemed to face insuperable difficulties, for there were apparently rigorous geophysical

¹⁹⁸ Kitcher (1990), 5

¹⁹⁹ Kitcher (1990), 6, my emphasis

demonstrations that the forces required to move the continents would be impossibly large." If, like the great majority of scientists of the time, one concluded – in apparent accordance with the best available evidence – "that the geophysical arguments really did expose the implausibility of Wegener's theory," then it was irrational to continue to defend it. Yet, some did, and eventually their doggedness was vindicated by the discovery of abduction and adduction zones and the development of plate tectonic theory, which today is universally subscribed to by geologists.²⁰⁰

It is once again important to stress that Kitcher does *not* endorse the view that the *individual scientists* who opposed Wegener's ideas acted irrationally or irresponsibly. On the contrary, he concedes that each individual geologist who denied the possibility of plate tectonics theory did so rationally, considering the bulk of the extant evidence then available. Kitcher's point is more subtle. He simply wishes to suggest that what may plausibly be called individually "irrational" behavior – in this case, Wegener's and his scattered followers' failure to update their beliefs in the face of a preponderance of evidence – may in due course wind up being beneficial to the epistemic community as a whole.

Though this point may be subtle, its full implications for the ways we think about individual epistemic behavior and the overall community's epistemic well-being are quite weighty. What Kitcher's argument reveals is that we may face a necessary tradeoff between (a) maximizing *each individual's* probability of adopting the correct belief in the here-and now and (b) maximizing *the collective's* probability of adopting true beliefs in the long term. I will call behaviors that serve goal (a) *individual* epistemic virtues and those that undermine it *individual* epistemic vices. Those that serve goal (b) I will call *collective* epistemic virtues and those that undermine it *collective* epistemic vices. The central point of Kitcher's paper is that sometimes behaviors that undermine goal (a) simultaneously serve goal (b). In other words, behaviors that may be legitimately dubbed *individual* epistemic vices can at the same time be *collective* epistemic virtues.

200 Kitcher (1990), 7-8

Example Two: Dogmatism and the Digestive Tract - Kevin Zollman's "The Epistemic Benefits of Transient Diversity"

The other work from philosophy of science I will use as an example, published two decades after Kitcher's, is Kevin J. Zollman's "The Epistemic Benefits of Transient Diversity." Like Kitcher's "The Cognitive Division of Labor," Zollman's essay begins with a vignette from the history of science, this one taken from the 20th century. The incident Zollman relates is the history of research on peptic ulcer disease (PUD). As Zollman explains, scientific researchers in the late 1800s developed two theses about the causes of PUD, both at the time supported by some evidence. The first thesis was that PUD was caused by bacteria. The second was that it was caused by excess acid buildup in the stomach. These two theses were subscribed to with rough parity until in 1954 "a prominent gastroenterologist, Palmer, published a study that appeared to demonstrate that no bacteria is capable of colonizing the human stomach" and "concluded that all previous observations of bacteria were a result of contamination." Since Palmer was a widely respected gastroenterologist whose study seemed to conclusively demonstrate the impossibility of bacterial survival in the stomach, and since the outdated equipment of early science was commonly subject to such elementary flaws, the result of his study was "the widespread abandonment of the bacterial hypothesis."²⁰¹ Nevertheless, a few rogue clinicians and scientists clung to the bacterial hypothesis, in some cases incurring censure for doing so. For instance, at least one practicing clinician in Greece was fined for treating patients with antibiotics and the work of researchers Robin Warren and Barry Marshall, who claimed to have observed *H. pylori* bacteria in a human stomach, was dismissed. So resistant, in fact, was the scientific community to the research of Warren and Marshall that, in a fit of desperation after not having their results taken seriously for a number of years, Marshall eventually drank a vial containing *H. pylori*, immediately becoming ill with ulcers and successfully treating himself with antibiotics. Eventually, Warren and Marshall were able to get their results taken seriously, after which they were replicated and confirmed. In 2005, they received the Nobel Prize in Physiology of Medicine for their efforts.

A superficial reading of the *H. Pylori* saga couched in terms of epistemic vice would attribute the scientific community's resistance to the bacterial hypothesis to a canonical epistemic

²⁰¹ Zollman (2010), 20

vice like dogmatism or hubris. But this is not the reading Zollman gives. On the contrary, he is highly sympathetic to the *individual* scientists who stood against Warren and Marshall. After all, he writes, each of these scientists "became aware of a convincing study, carefully done, that did not find bacteria in the stomach," against which stood the unsubstantiated and dubious-sounding claims of isolated clinicians and the occasional dogged researcher who must have looked to most scientists like an individual clinging too tightly to a favored hypothesis. Taking the former seriously and dismissing the latter seems under these conditions fully rational, and in fact, argues Zollman, "[h]ad the acid theory turned out to be true, the behavior of each individual scientist would have been laudable."²⁰² A further implication, which Zollman does not explicitly state but is, I think, just as apt, is that each individual case of a clinician continuing to treat PUD patients with antibiotics despite having knowledge of Palmer's study may well have been ethically *irresponsible* when judged in isolation using what the scientific community at the time considered the best-available evidence. After all, each of these decisions came at the risk of either directly harming patients or at least failing to administer the best-available treatment to them. Likewise, if we restrict our view to a single case, the perseverance of Warren and Marshall before they obtained sufficiently revealing results may for similar reasons be plausibly judged as irrational, since it went against the grain of the best-available evidence at the time, and in some degree morally irresponsible, since the preponderance of the scientifically agreed upon evidence strongly suggested research along these lines was bound to be a waste of time and resources and thus came at the opportunity cost of more promising lines of inquiry, which in turn could easily have led to additional suffering among ulcer patients for whom better treatments were not being developed as quickly as they might have been otherwise.

As he goes on to explain in the rest of the essay, in Zollman's view, the moral of the story of PUD is decidedly *not* a straightforward tale of Warren and Marshall as epistemic exemplars and the rest of the gastroenterology community as epistemic sinners. On the contrary, Zollman's aim seems to be that to demonstrate that some degree of one particular kind of epistemic behavior that can aptly be called a "vice" at the individual level – what Zollman calls "dogmatism," or excessive resistance to abandoning a belief in the face of contradictory evidence

202 Zollman 21. I disagree with Zollman on this particular point, but agree with the overall logic of his essay.

– may in some circumstances actually be a *virtue* when viewed from the perspective of the collective.

Three Contemporary Conversations

Let's now return to the present. As I mentioned earlier, the undesirable epistemic behaviors and habits – i.e., individual epistemic vices – of 21st century democratic citizens have prompted increasingly numerous commentators to warn of negative repercussions both for democracy and society at large. I mentioned three types of conversations in which this sort of argument is made with some frequency in section one. These examples by no means exhaust the possibilities. In fact, it does not seem like an exaggeration to say there is a general sense of epistemic crisis that runs through many of our political conversations in this, the “information age.” It can be spotted in the editorials that warn daily of the dangers of fake news, our growing apprehensions about political advertisements calibrated by big data to prey upon the epistemic weaknesses of citizens, and the oft-repeated lamentation that 21st citizens appear no longer to subscribe even to a common set of facts. In these conversations and more, the epistemic behaviors of individual citizens often tend to be condemned not merely as instances of *individual* moral or epistemic failure but as threats to *democracy itself*. This prevalent worry has moved numerous commentators to call for institutions both public and private to intervene in the communications sphere and stanch the bleeding, in some cases resulting in changes to the law.²⁰³ As we have seen, in combination with commentators' pervasive pessimism about the judgment of lay citizens, it also contributes to the adoption of a view of democratic epistemic well-being that heavily emphasizes the need for lay citizens to defer to the claims of experts and plays down the value of citizens relying on their own (deficient) powers of judgment.

Drawing upon the themes developed in the previous section, I now want to argue that those who voice the worries, many of whom subscribe to the views I have labeled “Progressive” in this dissertation, seldom if ever present sufficient evidence to substantiate the case that *even* the sorts of behaviors they legitimately name as individual epistemic vices ought to count as

203 In 2017, for example, Germany passed the Network Enforcement Act, or NetzDG, punishing information platforms for failing to proactively take down content deemed by the authorities “illegal,” and in 2018 the French legislature passed a measure authorizing judges to demand the removal of “fake news” during election cycles. In both cases, worries about the ongoing well-being of democracy were at the fore.

vices at the collective level. Commentators almost never trace a causal path from these behaviors to *collective-level* epistemic *outcomes*, nor even give any indication that there might be complications involved in scaling up from the individual to the collective level in this way. Instead, they appear to simply assume that the same behaviors that are detrimental to individual knowledge-seekers' tendency to find the truth are bound to be the same as those most beneficial to the growth of knowledge for the collective. I will argue that this assumption is a significant mistake. So significant is this mistake, in fact, that not we ought not only to hesitate to embrace the pessimistic conclusions about the state and direction of democracy so often connected to observations of individual-level epistemic vice, but in some important cases – indeed, some of the very cases most cited by Progressives as evidence of our collective epistemic “malaise” – we ought to see these very same behaviors as *beneficial* to democracy, or at least *indicative of* a democracy whose epistemic well-being is robust indeed.

I will make my case by responding to the three interrelated lines of conversation mentioned in Section One. These are (1) the conversation in the academy and the mainstream political sphere revolving around the role of "cognitive biases" and "motivated reasoning" in human psychology; (2) the conversation in those same arenas about *epistemic hubris*, or citizens' unprincipled reliance on their own judgment rather than that of experts; and (3) the worry that the ease and affordability of 21st century communications technologies are bound to drive citizens into ideologically isolated groups that foster political polarization, "extremism," and political ignorance.

Section Two: The Collective Epistemic Benefits of Motivated Reasoning and Other Cognitive Biases

In recent years, a number of political scientists have taken to studying the effects of "motivated reasoning" in the political arena. The basic idea behind the concept of motivated reasoning is that the beliefs people adopt and/or espouse are arrived at by a reasoning process whose aim is not *exclusively* to believe the truth. Instead, that process serves several aims, *one among which* – and not always the most influential one – is the desire to believe the truth. As Charles Taber and Milton Lodge put it, citizens' reasoning is motivated by both “*accuracy goals*,

which motivate them to seek out and carefully consider relevant evidence so as to reach a correct or otherwise best conclusion" as well as "*partisan goals*, which motivate them to apply their reasoning powers in defense of a prior, specific conclusion."²⁰⁴

William Minozzi and Michael Neblo have criticized Taber and Lodge, among others, for their work on motivated reasoning. However, their criticisms consist in contesting the *individual-level* irrationality of the behaviors named as such by "motivated reasoning" scholars; point out that as of yet we have no clear understanding of how pervasive these behaviors are; and arguing that there is hope that motivated reasoning may be "substantially remediated" through relatively simple interventions.²⁰⁵ While I think there is merit to these objections and that Minozzi and Neblo make a strong case that we ought to condition the readiness with which we label these behaviors individual epistemic vices, in this essay I will ignore these complexities and simply assume motivated reasoning scholars are largely correct in portraying motivated reasoning as an individual epistemic vice, and one that is more or less pervasive among the 21st century democratic citizenry. Even assuming all of this is true – even if motivated reasoning and other cognitive biases are detrimental to individual truth-seeking and are more or less rampant among the public – my argument is that this bears no necessary implications for the *overall* epistemic well-being of *democracy*. In fact, there is good reason to believe that some of the cognitive habits identified as cognitive biases are actually *beneficial* to collective knowledge-seeking.

To see how motivated reasoning could simultaneously be an individual epistemic vice yet be compatible with, or even contribute actively to, a community's overall epistemic health, let's revisit two of the historical examples above. Let us assume – not implausibly – that what inspired Joseph Priestley to cling so tightly to phlogiston theory at the turn of the 19th century was not a love for truth but rather *solely* his desire to vindicate his previous work and to build up his ego by defending the theory to which he had dedicated much time and effort and staked much of his reputation. Let us further assume – again, not implausibly – that similar motivations impelled Joseph Wegener and the others who along with him vociferously defended the theory of continental drift against what seemed at the time to be overwhelming evidence to the contrary. In

204 Taber and Lodge (2006), 756, their emphasis. Note that although Taber and Lodge offer "partisan" motivation as the sole other motivation that might influence citizens' reasoning, that term should not be read in the exclusively political sense as in "affinity for a certain political party." Instead, it may be read more generally as meaning "having an interest in the victory of one among various competing candidates."

205 Minozzi and Neblo (2020), 15-17

short, let us suppose their support of their own favored theories was heavily motivated by the sorts of biases contemporary social scientists have compellingly argued to characterize much of human thinking, *including* thinking in such seemingly non-political realms as mathematics²⁰⁶ and science.²⁰⁷ If, as may well have been the case, Priestley and Wegener were motivated by such base and epistemically unpraiseworthy aims, would the impact of their work on the legacy of chemistry and geology, respectively, be any different? I think the answer is obviously no. Whatever his motivations, Priestley's work served in the long run to strengthen Lavoisier's "new chemistry," both by forcing its proponents to come up with more and more compelling experiments and by providing a rival hypothesis victory against which, rightly enough, made the "new chemistry" look all that much more impressive. Indeed, to this very day, the work of Priestley and his fellow phlogistonians is taught to chemists in a way that demonstrates the explanatory superiority of our by now not-so-new chemistry. In this way, Priestley's challenge to Lavoisier not only failed to undermine the epistemic health of the scientific community but actively contributed to it. As for Wegener, both the scientific community and the rest of us get the benefits of his and his followers' dogged resistance to the consensus regardless of whether their motives in resisting stemmed from a high-minded fidelity to the truth or from a self-centered obsession with seeing their work vindicated. In both cases, then – one in which motivated reasoning happened to coincidentally align with the truth, and one in which it aligned with error – even if the *individuals* involved were guilty of epistemic vice, their work still served to strengthen the scientific *community* epistemically.

Now, let us consider a counterfactual world in which Priestley and Wegener had been free from the individual epistemic vice of motivated reasoning and instead been impelled entirely by the ostensibly purer motive of believing the “objectively” correct truth. Is it easy to say that the scientific community would have been better off in either case if so? I don't think so. It is not obvious that Priestley's conversion to the "new chemistry" – making him just one more scientist joining an already overwhelming consensus – would have helped demonstrate that theory's merit

206Nurse and Grant (2020)

207No serious scholar of the history, sociology or psychology of science I am aware of denies the presence of bias in the reasoning of scientists, and many think it is as or nearly as prevalent in such research as in everyday human thinking. See MacCoun (1998) for an extensive list of examples of bias in scientists' interpretations and uses of scientific research, as well as an extensive list of citations of similar examples of bias in the *conduct* of scientific research itself.

to anywhere near the same degree that his provision of a capable foil did. Turning to Wegener and his followers, would geology have been better off epistemically if their reasoning had been motivated only by the purest intentions? Again, it is far from clear that this is the case. On the contrary, a strong argument can be made – and indeed Kitcher makes the case – that the fields of chemistry and geology, respectively, were better off with Priestley and Wegener being motivated reasoners than they would have been otherwise. Nor does it matter that we cannot know whether *in fact* the reasoning of Priestley and/or Wegener was motivated in this way. Any sober surveyor of human affairs should agree with contemporary psychologists and political scientists that motivated reasoning is a pervasive part of humans' cognitive behavior, *including* that of scientists both past and present. Even if, as a matter of historical fact, Priestley and Wegener were exceptions and in reality were epistemic exemplars fully motivated by their love of truth, the scenarios laid out two paragraphs above are so fully consistent with what we ought to expect from humans in every walk of life that we ought to expect similar motivations to inspire a significant proportion of scientific work both past and present. We ought to expect many scientists to cling to pet theories longer than the evidence warrants – as, in fact, philosophers and sociologists of science since Thomas Kuhn have frequently argued to be the case.²⁰⁸ My claim here is that we ought not to *reflexively* lament that fact and assume that if only the reasoning of scientists entirely unmotivated the scientific community would be epistemically better off. For similar reasons, we should not assume that democracy would be better off if only citizens' reasoning was less motivated. On the contrary, I will later argue that the presence of at least *some* degree of motivated reasoning among the citizenry is likely to be good for the epistemic well-being of democracy.

Section Three: Anti-Vaccers as Epistemic Assets to Democracy

Another epistemic vice whose 21st-century iterations have given rise to a great deal of worry about the epistemic well-being of democracy is the vice of *epistemic hubris*, or groups of lay citizens' insistence on relying on their own judgment, rather than deferring to that of experts, when the two disagree. Lay citizens are routinely criticized, and often enough either explicitly or implicitly mocked, for relying on their own judgments against those of experts when the two

208 Kuhn (2012 [1962])

diverge. The disdainful light in which such behavior is viewed by Progressives is often difficult to miss. For example, Stephan Lewandowsky and colleagues imply that we are headed toward a "dystopian future" in large part because significant numbers of citizens prefer to base their beliefs on "an opinion market on Twitter" instead of the consensus of "97% of domain experts" on matters such as "whether a newly emergent strain of avian flu is really contagious to humans, or whether greenhouse gas emissions do in fact cause global warming."²⁰⁹ Elsewhere, James Traub encapsulated the interpretation of Brexit that prevailed, and continues to prevail, amongst many commentators when he characterized the vote as the victory of "mindlessly angry" citizens whose refusal to "defer to the near-universal opinion of experts" was equivalent to "denying reality."²¹⁰ Traub's characterization of citizens' refusal to "defer to the near-universal opinion of experts" as the democratically harmful equivalent of the cardinal epistemic vice, ignorance, is exactly analogous to the line of reasoning that commonly underlies the arguments of climate truthers who chastise citizens for refusing to defer to the frequently-invoked 97% of climate scientists who agree about human-caused climate change, or the even greater consensus that exists in the immunology community about the effectiveness and safety of vaccines. In all these cases, the *very act* of a lay citizen refusing to defer to the expert consensus is viewed as (apparently self-evidently) an instance of epistemic vice and, as such, is treated as a threat to democracy.

Now, I am fully in agreement that there are good reasons to trust the climate science and immunology communities on these and other issues, just as there are good reasons in innumerable other epistemic arenas relevant to 21st-century life for lay citizens to defer to the opinions of experts. I agree, then, that in many cases the refusal to defer to experts in these arenas is therefore a case of epistemic hubris, or excessive reliance on one's own judgment in a case where deference to someone better placed to render a good judgment would be more appropriate, which is rightly treated as an individual epistemic vice. Moreover, I grant that in the case of climate science specifically there is good reason to believe that the extent to which this vice reins among the public may well result in significant harms, and therefore ought to be

209 Lewandowsky et al (2017), 354. This tone of mockery is nothing new for Lewandowsky, who regularly subjects climate deniers to such treatment in a way that sharply contrasts with the advice he gives elsewhere about how to effectively persuade consumers of mis/disinformation (e.g. Lewandowsky 2019).

210 Traub (2016)

viewed as a *specific case in which* the proliferation of an epistemic vice among the public is likely to cause harm to the collective. However, this does *not* mean that as a *general rule* we can assume that a similar relation exists in all or even most cases. In fact, I will make the case that *collectively* and *in the long term*, even the anti-vaccination “movement,” such as it is, is both indicative of and contributes to the overall epistemic well-being of democracy.

Anti-vaccination groups have caught the attention of many commentators in the last few decades, with anti-vaccers commonly cited as examples of the proliferation of epistemic vice in this “post-truth” era and of the threat to greater society that this proliferation portends.²¹¹ However, I believe the best empirical evidence suggests these worries are dramatically overblown. Moreover, I believe that anti-vaccers raise some legitimate questions no one else seems to be asking about the ways vaccines are administered, the incentives involved in the manufacture and distribution of vaccines, and the legal structures that govern the immunization community. Since they seem to be persuading very few people to forgo vaccination, and since they are fulfilling a watchdog role over authorities no one else seems motivated to perform, I believe the anti-vaccination community is, on the whole, epistemically *beneficial* to our democracy, despite – or even *because of* – the proliferation of epistemic vice among the *individuals* comprising that community, which motivates their support for erroneous conclusions about the merits of vaccination and opposition to mandatory vaccination among the public.

A few paragraphs ago I hinted that I believe worries about the oft-touted “upsurge” in anti-vaccination movements to be dramatically overblown. This is because I take the primary – and best – reason to be concerned about people’s *beliefs* about vaccines to be that those who suspect them of being ineffective or dangerous will not receive them themselves or give them to their children, thus making those children vulnerable to vaccine-preventable disease and decreasing the security we all derive from “herd immunity.” I emphasize the need to focus on people’s *actions* because I believe a combination of common sense and a basic valuation of some level of autonomy and diversity dictate that people’s private beliefs ought not be considered public business unless and until they serve as a basis of action, and that our concern with those beliefs ought also to be scaled in rough proportion with the extent to which that action affects others (as best we can determine). The primary reason we ought to worry about people holding

211 E.g. Lewandowsky et al (2017)

and/or spreading the belief that vaccines are ineffective or dangerous, then, is if their doing so manifests in actions detrimental to the public good. Consistent with the arguments of the preceding chapter, this means that determining just *how* concerned we ought to be about anti-vaccers calls for empirical investigation. Just how much of an effect on people's actions *is* the so-called anti-vaccination "movement" having? Two empirical metrics would seem to be highly relevant to the question – the rate of vaccine coverage over time and the instance of vaccine-preventable disease. Has either of these increased during the heyday of those much-maligned platforms of social media and the blogosphere, so commonly blamed for serving as a permissive vector for misinformation and for enabling "extremists" to connect with like-minded people "with greater ease and frequency...without hearing contrary views?"²¹² Despite voluminous claims to the contrary, often forwarded by professional journalists²¹³ and in academic journals,²¹⁴ overall vaccine coverage rates have *not*, in fact, discernibly declined, either in the United States or globally in the era of the internet. This is true no matter whether we look at trends from the past decade or restrict our queries to the past few years. On the next page are graphs of the latest available data from the World Health Organization (WHO) and the United States' Center for Disease Control (CDC) tracking coverage for the standard battery of vaccines in the US over time. In no case is there ever more than a one-percent overall decrease in the proportion of children receiving a given vaccine, either within a given year or within the entire range of the dates covered by the data (11 years for the WHO data; 4 years for the CDC data), and in many cases the overall trend is an *increase* in coverage over time.

It is true that in 2018 the United States experienced the highest number of measles cases (1,282) since 1992. However, that is not indicative of a decline in *overall* vaccine coverage in 2018. Instead, "[m]ore than 73% of the cases were linked to recent outbreaks in New York."²¹⁵ particularly among certain sectors of the orthodox Jewish community, among whom in recent

212 Sunstein (2016), 77

213 E.g. PBS (2010)

214 For instance, Larson and colleagues claimed in 2011 that "current antivaccination groups have new levels of global reach and influence, empowered by the internet and social networking capacities" (Larson et al 2011). Elsewhere, Bean characterizes vaccine opposition as a "social movement, fueled in part by anti-vaccine information on the Internet," implying in this and similar passages that anti-vaccination movements are a *growing,* threat, even though elsewhere he admits that "[l]ow vaccination rates and resulting outbreaks of vaccine-preventable diseases have been shown to be geographically related to pockets of vaccine opposition," or in other words restricted to idiosyncratic communities (Bean 2011, 1874).

215 Patel et al (2020)

years a trend toward vaccine opposition has been noted.²¹⁶ The abnormally high number of observed measles cases in this case seems to be attributable to the dynamics of herd immunity, which makes it exponentially more likely for an individual who lives in a community with an abnormally high number of individuals who have not received a given vaccine – such as the tight-knit Orthodox communities of New York in 2019, the students of certain charter schools in San Diego in 2008,²¹⁷ and the citizens of Clark county in Washington in 2019²¹⁸ –to pass on the disease. While outbreaks like the one in New York receive considerable attention in the political commentary sphere, such attention obscures *overall* immunization trends, which have held rock-steady in the United States for decades, including the entire era during which the internet and social media have risen to prominence. Also during this period, the global vaccination rate for vaccine-preventable diseases has increased astronomically. Despite the urgent tone with which Progressive commentators tend to discuss the role of the internet and social media in facilitating the spread of mis/disinformation on vaccines, then, it is difficult to identify any clear detrimental effect such mis/disinformation, or the so-called "movement" responsible for spreading it, has on citizens' behavior *as a whole*. Indeed, the evidence suggests that whatever persuasive force the arguments of the anti-vaccers supposedly flooding the internet and social media have, it is consistently outweighed, both in the short and long term, by the average American's stubborn commitment to subjecting themselves and their children to the standard battery of vaccines endorsed by the immunology community. It would appear that the population of the kind of anti-vaccers we ought to be worried about – those who *take action* on the basis of their anti-vaccination point of view and persuade others to do so – is consistently small on the American scale and probably diminishing on a global scale even in the heyday of the internet and social media.

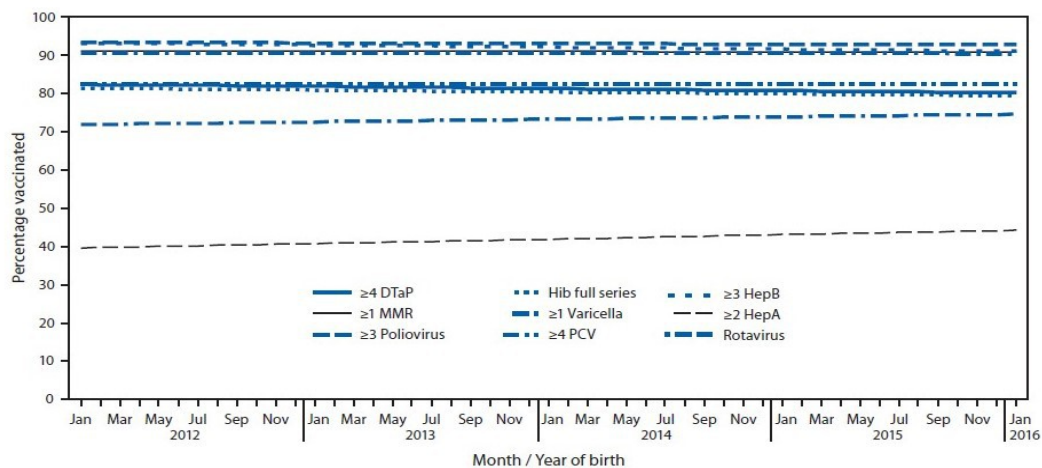
216 Feldman (2021)

217 PBS (2010)

218 Washington State Department of Health (2019)

Figure 1. Vaccine Coverage in the United States, 2012-2016²¹⁹

FIGURE. Estimated linear trend in coverage with selected vaccines* by age 24 months,[†] by month and year of birth[‡] — National Immunization Survey-Child, United States, 2013–2017



Abbreviations: CI = confidence interval; DTaP = diphtheria, tetanus toxoids, and acellular pertussis vaccine; HepA = hepatitis A vaccine; HepB = hepatitis B vaccine; Hib = *Haemophilus influenzae* type b conjugate vaccine; MMR = measles, mumps, and rubella vaccine; PCV = pneumococcal conjugate vaccine.

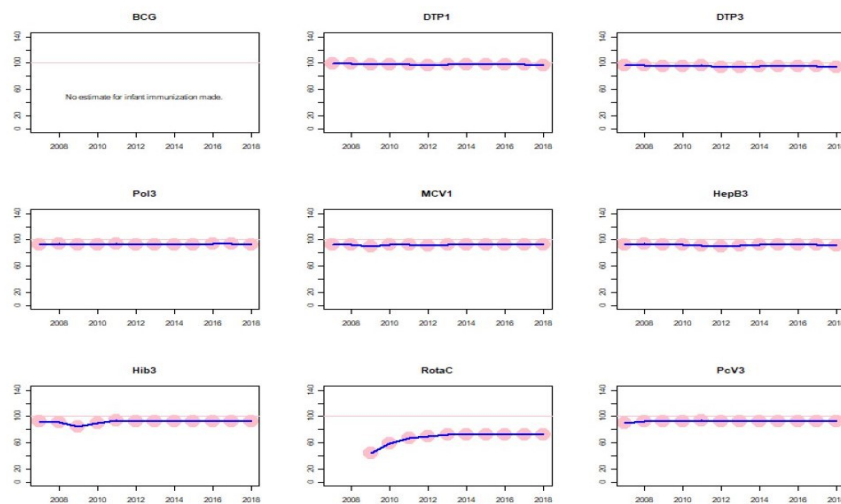
* Hib full series: ≥3 or ≥4 doses, depending on product type received (primary series and booster dose). Rotavirus: ≥2 or ≥3 doses, depending on product type received (≥2 doses for Rotarix [RV1] or ≥3 doses for RotaTeq [RV5]).

† Except for rotavirus, vaccination coverage was assessed before the child reached his/her 24-month birthday. The Kaplan-Meier method was used to account for censoring vaccination status for children assessed before age 24 months. Rotavirus vaccination was assessed before the child reached his/her 8-month birthday.

‡ Estimated linear relationship between month and year of birth and vaccination coverage, based on weighted linear regression analysis using the inverse of the estimated variance of each point estimate to construct the weights. Estimated percentage point change over 12 birth months: ≥4 DTaP -0.55 (95% CI = -1.20 to 0.10); ≥3 poliovirus -0.17 (-0.52 to 0.18); ≥1 MMR -0.11 (-0.58 to 0.35); Hib full series -0.51 (-1.13 to 0.11); ≥3 HepB -0.53 (-0.97 to -0.09); ≥1 varicella -0.05 (-0.53 to 0.42); ≥4 PCV 0.0 (-0.69 to 0.68); ≥2 HepA 1.13 (0.30 to 1.97); rotavirus 0.68 (-0.09 to 1.45).

Figure 2. Vaccine Coverage in the United States, 2007-2018²²⁰

United States of America: WHO and UNICEF estimates of immunization coverage: 2018 revision



July 2, 2019; page 1

WHO and UNICEF estimates of national immunization coverage - next revision available July 15, 2020

data received as of June 28, 2019

219 CDC (2018)
220 WHO (2018)

One might acknowledge these aggregate trends yet still contend that due to the *severity* of the threat posed by communicable illness to public health, even the apparently minimal impact of this small minority of anti-vaccers is harmful to the *overall, long-term* well-being of democracy. I disagree. In fact, I believe just the opposite. Recall the history of vaccine coverage and anti-vaccination movements related in Chapter Two. Contrary to what is implied when commentators attribute anti-vaccination sentiment to digital technologies various anti-vaccination "movements" have arisen periodically ever since the 1950s, which is when most developed countries began administering roughly the current standard battery of immunizations. While these movements have always aroused the alarm of public health officials and scholars in their day, they have rarely caused a significant dent in overall vaccination levels.²²¹ Moreover, in each of the eras during which some causal role has been plausibly tracked from anti-vaccination movements to declines in vaccine coverage – most prominently, the decline of pertussis coverage in Sweden, England and Australia in the 1980s-90s – it was not the absorption of low-quality information drudged up from the tabloid stand or any other non-professional or non-expert informational source that started the "movement." Instead, it has typically been opposition to the vaccines arising from *within the scientific and public health communities*.²²² Indeed, the most famous vein of vaccine opposition – that based on the erroneous idea that the measles, mumps and rubella vaccine causes autism – was initiated by a paper published by a practicing scientist in the highly prestigious medical journal, *The Lancet*. Despite occasional examples like this, where "rogue" scientists and/or practitioners have come out against vaccines, democratic citizens have shown a consistent historical propensity to embrace vaccines at ever-increasing rates, with only a few short-term blips noticeable in vaccine coverage even at the height of the "movements" these figures inspire.²²³ Along with contemporary empirical trends, this historical record suggests that lay citizens' powers of judgment are sufficient to allow the public to routinely fend off anti-vaccination movements' persuasive advances and, in tandem with moderate incentives such as those currently employed by various U.S. governments, to choose to have their children

221 Gangarosa et al (1998), Blume (2006), Vanderslott et al (2019)

222 Gangarosa et al (1998)

223 It is worth noting that the instance of pertussis worldwide has resurged somewhat in the last several decades. However, this resurgence is *not* due to a decline in coverage but rather to a combination of (a) the waning efficacy of the vaccine within an individual's lifespan, (b) lower overall effectiveness of acellular pertussis vaccine, which produces an immunological response that is slightly different than that produced by pertussis itself, and (c) genetic variations that make vaccines less effective against certain strains of *B. pertussis* (Esposito et al 2019).

vaccinated. Because the vast majority of citizens have consistently shown that they are capable of using their judgment to choose the right action with regard to vaccination, I argue that we have little reason to fear many of them being coaxed into doing otherwise by whatever anti-vaccination "movement" exists in the blogosphere and on social media platforms poses any threat either to public health or to the epistemic well-being of democracy. Because of this, I believe there is no good reason to suppose anti-vaccers armed with 21st century communications technologies pose a significant threat to the epistemic well-being of democracy.

But I will not rest my argument there. Instead, I want to argue further that anti-vaccers are both a *sign of* and a *contributor to* the epistemic well-being of our democracy as it stands. To see why, I will do something anti-vaccers' critics typically do not do. I will look at the actual arguments the standard-bearers of this so-called "movement" tend to muster in defense of their opposition to vaccines. Three of the most common reasons cited by anti-vaccers for their opposition to mandatory vaccines are that (1) vaccines cause autism, (2) vaccines produce side-effects that are themselves more threatening to the health of children than the (often rare) diseases the vaccines are intended to prevent, and (3) vaccines are part of some kind of conspiracy, either one perpetrated by governments and pharmaceutical companies to make money²²⁴ or, more sinisterly, to sterilize a given populace.²²⁵ Though all three lines of argument are commonly rebuked, if not openly mocked, by vaccination proponents and held up as signs of the sort of "post-truth" era individual vice that spells trouble for democracy, I believe the more appropriate takeaway is just the opposite: by voicing these concerns in the public sphere, anti-vaccers are performing the vital societal role of maintaining *vigilance toward authorities* in matters toward the rest of the public is in real danger of becoming complacent.

Let me explain my reasoning by taking anti-vaccer objections (1), (2), and (3) in turn. As mentioned above, the commonly subscribed-to idea that vaccines cause autism originated not on Twitter, Facebook or some obscure, politicized blog, but rather in the pages of the prestigious medical journal, *The Lancet*, in an article published by Andrew Wakefield and a dozen other colleagues in 1998. Defenders of the scientific community tend to view the eventual retraction of that article by *The Lancet* as evidence of "science policing science," but that is a misleading

224 Sales (2019)

225 Kenya Conference of Catholic Bishops (2014)

characterization. In fact, when controversy arose after the article's publication *The Lancet* initially "exonerated Wakefield and his colleagues from charges of ethical violations and scientific misconduct," despite it having been revealed that "Wakefield had been funded by lawyers who had been engaged by parents in lawsuits against vaccine-producing companies." Finally, in 2010, fully twelve years after the publication of the original article and a decade after it had been notified of its problems, *The Lancet* published a full retraction "as a small, anonymous paragraph." Even *that* move, however, was not prompted by an outcry among epidemiologists. Instead, "the exposé" that led to the retraction "was a result of journalistic investigation,"²²⁶ specifically the work of Brian Deer of the *British Medical Journal*,²²⁷ who was not a trained immunologist but instead a journalist, i.e. a lay person with regard to epidemiology. Though those who oppose vaccines on the grounds that they cause autism are often held up as examples of what is wrong with lay citizens choosing to rely on their own judgment rather than that of the scientific community, or what I have called *epistemic hubris*, the moral of the *Lancet* saga may arguably be just the opposite. That is, the moral of the *full* story of the *Lancet* saga may be to demonstrate the *value* of lay citizens being willing to question the scientific community, and the risks of excessive deference to it, as the reverse.

Argument (2) among those I mentioned as commonly cited by anti-vaccers was that some vaccine regimens may entail risks to children that outweigh the threat posed by the diseases the vaccines are intended to prevent. A frequently aired version of this worry revolves around mercury, which is used as a preservative in a number of vaccines. Because many people know that mercury can be toxic, this has raised concerns about the threat posed by vaccines themselves to small children. This concern, however, is misguided, as the type of mercury included in vaccines, ethylmercury, is cleared from the body quickly and easily, making it very different – and much safer – than the methylmercury found, for example, in some kinds of fish. Moreover, ethylmercury is not in the vaccinations given to small children in the United States, and has not been since 2001, though it is still a component in some vaccines typically administered at an older age.²²⁸

226 Sathyanarayana and Andrade (2011)

227 Deer (2007), Deer (2011a), Deer (2011b)

228 CDC (2020)

Insofar as their objections revolve around mercury *specifically*, then, vaccination opposition seems clearly to be misguided. However, is the *general* worry that the standard vaccination regime may be predicated on an underestimation of the negative side-effects of vaccines unreasonable? I think not. As advanced as medical science is today, the human body is an enormously complicated and dynamic thing, and physical health the result of interactions between genetics and environment whose complexity far outstrips our current understanding. The 20th century is replete with examples of medical treatments whose side-effects, in retrospect, we judge to far outstrip their benefits. Examples include lobotomies, shock therapy, quaaludes, and diet pills like Fen-Phen. Arguably, the same could be said in relation to the current opioid crisis, which many view as the worst addiction crisis in United States history, and whose origins and scope owe at least partially to the medical industry's acceptance of pharmaceutical companies' assurances that patients would not become addicted to the powerful opiates they were producing.²²⁹

The examples just mentioned are commonly acknowledged as failures on the part of the medical community, but I want to point out a difference between the approach I am advocating here and that which seems to be the preference of outraged political commentators. It is common for the latter to lambast the medical community for historical failings like the ease with which they approved prescriptions for opioids, diet pills, etc., in *moral* terms, as if they are primarily examples of corruption or a bad culture among the relevant medical practitioners. While I certainly do not wish to deny that such corruption is always a danger, my inclination in cases like those mentioned above is to look *first* to simple human fallibility as the culprit. The medical industry is no more omniscient than any other group of human beings has ever been. Their training, as indispensable as it is, comes with the same risks of inurement to convention and what I call *monoculism*, or the tendency to view something from a unitary point of view, as all other forms of intensive education in a given discipline. There is also a certain status attached to being a medical researcher or practitioner that can easily lead those who have attained it to be overly dismissive of the arguments made by those who have not, a tendency no doubt only further encouraged by the exhaustion many such experts must deal with by being pelted with objections, questions and comments by lay persons that are legitimately ill-founded. Try as we might to

229 HHS (2020)

mitigate these dangers by subjecting medical treatments to empirical tests, such tests are necessarily limited. We cannot expect scientists to always secure perfect knowledge of a medical treatment's average benefits vis-a-vis its average risks until a large, diverse sample of people undergoes that treatment for a long time, nor should we expect scientists to be wholly immune from the dangers of hubris and exhaustion that are often part and parcel of attaining the status of an expert. Acknowledging this, I believe we ought to approach failures of expert communities such as those mentioned above with a certain predisposition toward charity.

At the same time, however, we should not ignore the fact that doctors and medical researchers, like all expert communities, often face perverse incentives when it comes to their willingness to subject treatments they invent and/or prescribe – incentives that include both monetary and reputational stakes, as well as what Stephen P. Turner has aptly named as the eternal and ineliminable interest of all expert groups, the desire to continue to be seen as an expert, i.e., someone who "knows something that the consumer of expert knowledge does not and cannot know."²³⁰ Because of the immunology community's fallibility and because they, like all other experts, sometimes face perverse incentives, there will always be a need for people who do not belong to that community – in other words, non-experts in the realm of immunology, or *lay persons* – to keep watch over them, as illustrated in the case of *The Lancet* and its resistance to retracting the article that attributed autism to the MMR vaccine. In many cases such watch-keeping by lay persons will only serve to affirm that the experts being watched over do in fact deserve to be trusted, as I take typically to be the case with the medical community. But the very propensity for a community to produce consistently reliable knowledge, and good behavioral subscriptions on that knowledge's basis, is likely to encourage complacency in the public over time. The more reliable a community's claims prove to be over time, the more automatically the lay public is likely to be to simply take them at their word. A common sense understanding of human character and motivations strongly suggests that this is likely to invite both complacency and corruption in that community whose trustworthiness is taken for granted. Thus, it is precisely in those instances where experts have proven consistently trustworthy that the claims of the most eccentric and outlandish-seeming dissenters should be most jealously protected. This is both because the commonness with which the consensus view is accepted is bound to make *all* dissent

230 Turner (2003), 12

seem eccentric and outlandish, as well as because the dissent so indispensable to the growth and preservation of human knowledge is in such cases only likely to survive in communities whose beliefs seem to the rest of us eccentric and outlandish. If a community of experts enjoys such a lofty reputation of credibility and honesty that nearly all its claims are accepted with little question by the public, any mistakes, laziness or corruption on the part of those experts is most likely to be discovered by those few eccentric and outlandish individuals who have seen fit to contest the claims and authority of those experts.

Because of all this, I believe from a collective epistemic perspective it is more productive to view anti-vaccers who object to vaccines on the grounds of objection (2) as valuable assets to, and indicators of, our collective epistemic well-being, rather than threats to it. They are *valuable to* democracy because they keep dissent and vigilance alive among the public, which serves to counteract the public's natural tendency toward complacency toward authorities whose legitimacy is – in this case deservedly, in some others not – widely accepted by the public. And they are *indicators of* democracy's epistemic well-being both because their *presence* reassures us that our current informational environment encourages such communities of dissent to arise and because their *scarcity* and *lack of influence* reassures us that most of the public's judgment is sufficient to continue to believe in the value of vaccines even though they have access to the arguments of anti-vaccers.

Let us move on to argument (3), which sees vaccines as a conspiracy either to rob the public of its money or to sterilize some sub-population of "undesirables." By far the most common response to this line of argument is to laugh it off as patently absurd. In this, as in all other cases of laughing off "conspiracy theories" *merely* on the basis of their *being* conspiracy theories, i.e. accounts of the secret machinations of those in power most of us find implausible, I believe such laughing-off is *itself* a form of epistemic hubris.²³¹ In the case of vaccines in particular, it is also a sign of a poor historical and psychological understanding. Contrary to what anti-vaccers' critics often suppose, not only is the idea of a conspiracy between public officials and health care providers either to make a profit or to covertly sterilize some group of societal "undesirables" well in line with common-sense understandings of the types of human motivations that often prevail where money, power and racial-ethnic superiority are concerned;

231For an in-depth defense of a position compatible with my own, see Coady (2011, ch. 5)

they also have extensive historical precedent. Profiteering by health care providers through the prescription and provision of phony or ineffective medicines would certainly be nothing new.²³² Nor would covert sterilization done under the guise of public health, which, shocking though it may be to those with a delicate worldview or an Enlightenment-esque ruddy optimism about how much more "civilized" our era is than those that preceded it, has a long, wide-reaching, and in some cases quite recent history.²³³ Though I myself strongly doubt that the vaccination programs currently offered in the United States and most other countries at present are conspiracies of this sort, like most citizens I cannot personally oversee the manufacturing process of the vaccines or know with any high degree of clarity the intentions of those in charge of forming and overseeing it. Instead, my opinion is necessarily based on a series of extrapolations from my views about human nature, the effectiveness of gatekeeping mechanisms in the medical profession, the utility of "muckraking" efforts in a country with freedoms of speech and press, and dozens other of my general impressions about the way things are. I am fairly confident in my opinion that the vaccination programs encouraged by the United States medical community are not part of a conspiracy – confident enough to support their being made mandatory by, for example, requiring school children to receive them before being allowed to enroll. However, given its basis on a series of *general impressions* about the *way things are*, my confidence in the conclusion that vaccines are not a conspiracy against the public should *not* be so great that I should feel free to *dismiss out of hand the very idea* that political officials might conspire with health care providers to sinister ends, *especially* since the willingness to do so, either for profit or to control some despised sector of society, is compatible not only with common-sense notions of human psychology but also with extensive historical precedent. Responding to such claims by laughing them off cannot, then, be considered a responsible epistemic move and is, in fact, only likely to

232 Indeed, considering the unprecedented costs of health care the incentives to engage in this sort of behavior have evidently never been higher, so we should not be too surprised to learn about instances of such fraud being committed in our day - as, indeed, they are, with what justice department "authorities described as one of the largest health care frauds in United States history," which included the provision of unnecessary treatments at exorbitant prices, being prosecuted as recently as last year (Chokai and Jacobs 2019).

233 In the United States alone, programs have been carried out to sterilize members of undesirable racial, socioeconomic and cognitive status, all of which were at least partially accompanied by deception (Lawrence 2000, Andrews 2017). Worldwide, numerous similar examples can easily be found, many from recent history (e.g. McDonough, 2013, Bi and Klusty 2015).

make the sorts of abuses feared by anti-vaccers more likely in the future by reassuring those who would wish to carry such programs out that the public will never suspect them of doing so.

The point I want to press with most force, however, is that even if the willingness to invoke objection (3) as a reason for opposing vaccines *is* the result of epistemic vice on the part of most or even all of those who do so, the current anti-vaccination "movement," such as it is, is *still* both an *indicator of* and an *active contributor to* the epistemic well-being of our society and not, as Progressives take it to be, an indicator of or contributor to its deterioration. Recall the statistics reported several paragraphs above, which indicated that overall vaccine coverage has not displayed any discernible decline in the past few decades and has in fact, on the global scale, increased dramatically. Now, consider the stakes entailed in a parent having their child vaccinated, which, according to the most commonly-forwarded theories, include either an increased risk of autism, poisoning or sterilization. These stakes seem sufficiently consequential that one might reasonably expect that the vast majority of parents who *actually* believed the arguments of anti-vaccers would refuse to have their child subjected to such treatment, even in the face of legal deterrents such as not having their child admitted to public school. It is hard to believe that many parents would choose to forcibly sterilize their child or expose them to dramatic threats to their mental and physical health rather than to pursue alternative educational routes like home schooling groups or private schools if they *truly* believed vaccines to be secret sterilization programs or come with a significant risk of causing autism. In this case, then, the ongoing willingness of parents to have their children vaccinated – their "revealed preference" – seems likely to be a better indicator of their true belief, or at least the belief which they are willing to act upon (which in my view is the same thing), than whatever they may be posting on their blog or linking to on their Facebook home pages and Twitter feeds. If this is so, it follows that the number of actual, actionable subscribers to anti-vaccination theory – and certainly the number of people who subscribe to a degree sufficient to subsidize action – is small.

On top of this, consider the above-mentioned history, spanning across generations and across continents, of conspiracies to either sell useless/harmful treatments to the public under the guise of "medicine" or to purposely undertake sterilization programs of "undesirables" under that same guise. Consider, too, the easy compatibility of such behavior with motivations we know to be common in human beings – notions of racial/cultural superiority, disdain and disgust for

marginalized sub-groups within a given society, and the conscience-numbing allure of profits, the last of which has arguably never been higher, given the astronomical and still-rising costs of health care in the present day. And consider that neither scientists nor any other group of epistemic authorities seems to be any more immune to these motivations than the rest of humanity. Holding all this in mind, I ask: Is even a society in which conspiracies like these are not, in fact, *currently* taking place epistemically better off if its entire populace is fully trusting of and quiescent to the consensus of the immunology community and the claims of public health officials? Or would that society, instead, be better off in the long run if some small subset of its population stubbornly refused to accede to that consensus, even in the face of what ought from an individual epistemological perspective to be persuasive reasons to accede?

I argue the latter. My reasoning can be explained in terms of a straightforward risk-reward analysis. The threat of the public, or at least vulnerable sub-groups within the public, being taken advantage of by some combination of public officials and health care providers is well established and apparently chronic. As pointed out above, it spans multiple generations and every continent. Meanwhile, despite the notion's apparent currency among a substantial subset of worried political commentators, there is little if any evidence to support the idea that anti-vaccination movements ever have posed or are likely to pose any substantial threat to the public's well-being, epistemic or otherwise.²³⁴ That is to say, the rate of infection by vaccine-preventable disease in countries in which the vaccine is readily available continues to be vanishingly small, both in comparison to the total population and to the baseline rate of infection we should, on historical evidence, expect if the rate of vaccine coverage was not high. Though "outbreaks" do occasionally occur, they have, ever since the 1970s, always been both small in magnitude and isolated in spread. They amount to a stochastic distribution of temporary upticks in the infection rates of small, isolated and otherwise idiosyncratic communities.

I believe the proper conclusion to reach once one balances these considerations is that citizens have good reason to suppose that there is a *standing danger* of some meaningful number

234 To press this point home just once more, the largest so-called "outbreak" of measles in the past quarter-century consisted in 1,282 cases total in a single year, more than 800 of which were confined to a single community. Compare this number to the national average of nearly 550,000 cases/year through the 1950s (CDC 2018). The most noteworthy decline in vaccine coverage for this disease, in other words, only succeeded in temporarily boosting it to less than 0.2% of its *average* level among the population in the pre-vaccine era - an era during which the overall population of the United States was half what it is today.

of public officials and health care providers abusing the public if they can get away with it while they have little reason to suppose that either anti-vaccination movements or mis/disinformation about vaccines spread online *ever have* posed or *are likely in the future* to pose a significant threat to public health, or even to dissuade large segments of the public from getting themselves and their children vaccinated. Contrary to the democracy-threatening persuasive force Progressives often portray them as being, the anti-vaccination community that so stubbornly questions the medical community's consensus on the effectiveness of and risks posed by vaccines is better described as a small group whose most potentially dangerous arguments appear not to ever have been taken very seriously – that is, not seriously enough to *act upon* – by the overwhelming majority of the public. Since so little of the public seems to take their arguments seriously, anti-vaccers do not appear to be any sort of threat to that public. However, anti-vaccers' paranoia is very likely to motivate them to be abnormally vigilant with regard to the practices of the medical community and its possible collusion with officials – a vigilance which, I have argued, is both (a) always going to be necessary in order to guard against abuse by members of those communities and (b) likely to naturally recede amongst the bulk of the lay public in precisely those circumstances in which a given community has proven consistently reliable. As time goes on, we should expect complacency to grow amongst the general public with regard to such watch-keeping, so that they ask few probing questions of the most-trusted expert communities, such as immunologists and the medical community as a whole. Under such circumstances, we should also expect dissent to linger on only amongst those who are abnormally dogmatic and otherwise deviant. Since high levels of dogmatism and extremism are often enough the product of epistemic vice, we ought to expect such communities to manufacture and keep alive plenty of bad ideas. However, we ought also to expect that over time these communities will increasingly be the *most* likely among those in a given society to discover and keep alive legitimate and well-founded objections to the expert consensus, if only for the reason that they are the *only ones* who see any reason to subject that consensus to much scrutiny. Indeed, I believe this to be precisely the case at present in the anti-vaccination community, many of whose objections are ill-founded but a few of which – for instance, the question of why

vaccine producers have full immunity from liability suits²³⁵ under U.S. law²³⁶ – appear to be perfectly legitimate issues for public discussion, yet are by all indications seldom if ever even contemplated as possibilities by Progressives.²³⁷

By keeping alive legitimate questions surrounding the way vaccines are administered and regulated, anti-vaccers are performing an important epistemic role for society. Their motives for performing that role may well be irrational and their thought process excessively dogmatic, paranoid, or infected by one or more of any other number of *individual* epistemic vices, but none of that prevents them from playing a salutary role in maintaining and actively contributing to the epistemic well-being of democracy *as a whole*. In other words, however much their opposition to vaccines may be the result of individual epistemic vice, I take their presence as an abnormally vigilant subset of the community in an arena that needs constant vigilance to be a *collective* epistemic virtue. Since the anti-vaccination community appears to pose very little threat to the epistemic well-being of democracy by way of persuading the public not to immunize themselves or their children at the same time as it performs a valuable function of vigilance, I believe that their presence in 21st-century America actively contributes *to* our collective epistemic well-being and does little to nothing to detract from it. Our democracy is, in short, better off epistemically *with* pockets of epistemically vice-ridden anti-vaccers than it would be without them.

Section Four: The Collective Epistemic Benefits of Echo Chambers and Filter Bubbles

The third conversation in which the epistemic vices of lay citizens are blamed for bad collective epistemic outcomes is the series of discussions surrounding "echo chambers" and "filter bubbles." First, let me define these terms. The term "echo chamber" is a metaphor for a

235 A Voice for Choice (2020)

236 U.S. Code section 300aa-22

237 Certainly, they are seldom met with an open ear by anti-vaccination opponents, who met with such reasonable-seeming arguments generally refuse to acknowledge them, showing instead a strong preference for endlessly reiterating the basics of vaccine theory, often accompanied by bromides about the importance of trusting scientific authorities. For instance, of the first ten Google search results I performed on on the term "vaccine opposition" – all of which links were, incidentally, directed *against* anti-vaccers and in favor of the mainstream scientific consensus – only one mentioned any issue of legal liability. To its credit, *The New Atlantis* called this a "legitimate concern" and went on to explain why other measures taken by the U.S. government ought to mitigate it in the eyes of the public (Rothstein 2015). This was, however, the sole exception, despite liability being the *number one objection* to current vaccination policy listed by prominent anti-vaccination group A Voice for Choice (A Voice for Choice 2020).

self-selected community of like-minded individuals who are exposed only to arguments from co-partisans who "echo" their own beliefs back to them. The term "filter bubble" is a metaphor for the various ways in which an individual's conscious and unconscious choices serve to limit their exposure to information in ways likely to bias their worldview. The most prominent champion of the view that 21st-century communications technologies are damaging democracy by causing echo chambers to proliferate is Cass Sunstein, who argues that echo chambers serve to foster ideological extremism, polarization, and political ignorance.²³⁸ Meanwhile, the term "filter bubble" was invented by Eli Pariser, who worries that the prominence of choice-driven algorithms, combined with the profit motives of major internet companies like Facebook and Google, is likely to limit users' exposure to information that are bound to "skew [their] perception of the world,"²³⁹ thereby impeding us from becoming aware of and reaching agreement about how to solve important political problems. These arguments tracing the advent of increased choice, low-cost, and high-ease brought about by 21st century developments in communications technologies to negative political outcomes are integral to the Progressive position and are immensely popular in both the academy and the mainstream, with many commentators arguing or strongly implying that they constitute a threat to democracy itself.²⁴⁰

Now, it is first of all worth mentioning that, as was the case with Progressive depictions of the number and influence of anti-vaccers, there is good reason to think that the prevalence of echo chambers and filter bubbles and the role played by the technologies most commonly blamed for them – namely, social media platforms like Facebook and Twitter – is overstated, perhaps dramatically so. In fact, the empirics suggest that in many ways the truth may be the *exact opposite* of what the echo chamber/filter bubble narrative supposes – that is, on average people may actually be exposed to a *greater* diversity of worldviews through social media and the internet more generally than they would be without them.²⁴¹ My point in this sub-section,

238 Sunstein (2016)

239 Pariser (2011), 14

240 For just a few examples, see Prior (2009), El-Bermawy (2016), Deb, Donahue, and Glaisyer (2017), and Schiffrin (2017).

241 The evidence that this is the case is quickly becoming voluminous. For example, Gentzkow and Shapiro found that people who consume news online are exposed to significantly more diverse ideological viewpoints than is characteristic of their "offline" communities (Gentzkow and Shapiro 2011), Flaxman and colleagues found that social media use is "associated with an increase in an individual's exposure to material from his or her less preferred side of the political spectrum" (Flaxman et al 2016, 298), and Dubois and Frank found that exposure to diverse ideological points of view was positively correlated with interest in politics among internet users (Dubois and Blank

however, is not to simply report the growing consensus that the *actual* effect of 21st-century communications technologies appears to be, if anything, to *reduce* the prevalence of echo chambers and filter bubbles in our society, but to contest the near-ubiquitous assumption that a reduction in echo chambers and filter bubbles would necessarily *good* for the overall epistemic well-being of democracy, an assumption that is nearly always accepted without any attempt at justification.²⁴² Since its proponents so seldom attempt to explain why they assume ideological isolation, ideological reinforcement and the growth of “extremism” among *individual* citizens are bound to have negative effects on democracy *as a whole*, it is impossible to be sure what their reasoning is, but it seems to me that the most plausible reason is that they assume that what rightly counts as an *individual* epistemic vice – in this case, isolating oneself from alternative viewpoints and basking in the exaggerated confidence that comes from having one's beliefs affirmed by those who share them – is bound to produce negative effects at the level of *collective* epistemology. As the reader will no doubt have guessed by now, I believe this assumption to be deeply misguided. Not only do I believe these individual vices do not always entail deleterious collective outcomes; I also believe their presence in a subset – perhaps even a substantial subset – of individuals in a given community may actually produce epistemic benefits to the collective. To make my case, I will cite two lines of argument, one taken from the Zollman article introduced in Section Two of this chapter and one encapsulated in an article written by legal scholar Heather Gerken explaining the concept of “second-order diversity.”

Let's begin with Zollman's article on the epistemic benefits of “transient diversity.” The first part of Zollman's essay is dedicated to relaying the history of peptic ulcer research, which I reported above. However, the greater part of that essay is not historical but formal. In it, Zollman uses computer simulations that model the human knowledge-seeking process in hopes that those

2018, 740). Because of results like these, political communications scholar Kelly Garrett gives voice to the growing consensus among social scientists studying internet users' information diets that “echo chambers are not a typical part of Internet users' experience” (Garrett 2017, 370).

242 See, for instance, UNESCO's assertion in its Guide to Fake News and Disinformation that by “[reducing] many individual users' exposure to alternative views and verified information” the advent of social media “has amplified the risks associated with ‘information disorder’” (UNESCO 2018, 62). It is worth mentioning that, like many others who casually invoke the term, UNESCO never mentions what it means for something to count as being “verified.” The term “information disorder” refers to the framework of Christine Wardle and Hossein Derakhshan, which has become popular among those studying the epistemic impact of 21st-century communications technologies on democracy. In their explication of that framework, Wardle and Derakhshan make the same automatic assumption that, to the extent they exist and foster polarization, political extremism, ideological isolation, echo chambers and filter bubbles are worrisome for democracy (Wardle and Derakhshan 2017, 51).

simulations can provide insights into the way(s) individual-level behaviors affect collective epistemic outcomes. In the simulations, the knowledge-seeking process is modeled as a "bandit problem," a form of dilemma discussed by decision theorists. The key features of the "bandit problem" can be illustrated through the example of someone given a choice between trying their luck on several different slot machines with no *a priori* knowledge of the odds of the machines' payouts. In that situation, the gambler must make a choice about how to allocate their limited number of coins without actually knowing which machine is most likely to maximize their payout. This situation, argues Zollman, is analogous to that faced by researchers such as those studying peptic ulcer disease before the discoveries of Warren and Marshall. As Zollman says, during this time researchers needed to choose whether to "dedicate their time to developing more sophisticated acid reduction techniques" or to use their limited resources in time and energy to "search for the bacteria that might cause PUD." Because what most scientists viewed as the best-available information at the time – the results of Palmer's study – seemed to strongly suggest that bacteria *couldn't* exist in the human stomach, the scientific community viewed the bacterial hypothesis as a waste of time, and a potentially a harmful one, since every moment spent meandering down hopeless avenues of research was a moment that could have been used to employ most plausible research to date to ease the pain of people suffering with ulcers.²⁴³ Zollman argues that many scientific problems are like the "bandit problem" and the case of PUD research. That is, they require researchers to allocate finite resources without knowing the actual odds that they are following a line of inquiry that will lead to an improvement of knowledge and not meandering down an epistemic road to nowhere, and the stakes of their decisions entail serious risks.

I believe that not only was Zollman right to conceive of scientific inquiry in this way, but that this conception applies to *all* human inquiry, though the scope of the risks involved in choosing which among multiple alternative hypotheses to pursue varies from case to case. My reasoning derives straightforwardly from the assumption of chronic epistemic fallibility, i.e. the assumption that any particular belief we currently hold might either be *wrong* or, in terms I usually prefer, be *inferior* to some better explanation that meaningfully diverges from the one we currently endorse. Chronic epistemic fallibility means that we simply cannot tell, for any

243 Zollman (2010) 22-7

particular belief, whether or not some new information or perspective will come to light which shows our current belief to be either in error or at least explanatorily sub-optimal. It is far from uncommon to discover that beliefs long subscribed to by individuals, groups, and even entire epistemic communities – such as the community of gastroenterologists – were in error or sub-optimal. Moreover, it is common for the idea that reveals the flaw in our previous understanding to come from a surprising and unanticipated avenue of inquiry. Often, this comes in the form of a fluke accident, such as when Alexander Fleming discovered penicillin by forgetting to wash his storage vessels or when Percy Spencer's decision to bring a chocolate bar to work one day resulted in his discovery of the microwave when it unexpectedly melted in his pocket. No less common a source of discovery is for insights from one discipline to yield unanticipated value for another, such as Herbert Simon's work in economics, which resulted in the "bounded rationality" model of human decision-making that has massively impacted psychology and all the other social sciences. The most famous scientist of the 20th century and perhaps all time, Albert Einstein, may plausibly be described as an example of both these phenomena - discovery through serendipity and discovery through unanticipated extra-disciplinary insight. Einstein was, of course, a long-time patent office worker whose post happens to have been located in Bern, Switzerland, in an era during which numerous innovations involving the coordination and calibration of time-keeping implements just happened to be passing through that office.²⁴⁴ These and countless other examples taken from the history of human intellectual progress should make it abundantly clear that we cannot estimate to any great degree of certainty the odds of any particular avenue of inquiry resulting in the discovery of useful knowledge, nor which problems are best addressed by the insights of whatever specific line of inquiry we pursue. The conditions of human knowledge pursuit are, in other words, the epistemic conditions described in the "bandit problem."

Zollman's exploration of the implications of the "bandit problem" is of particular interest in this section of the chapter because of its implications for the collective epistemic effects of ideological isolation. In Zollman's simulations, 10,000 "agents" were programmed to search for the optimal solution to some problem. In this case, the quest for that optimal solution was

²⁴⁴ Peter Galison has persuasively argued that precisely this unplanned and frustrating relegation to patent-office obscurity may well have served as the impetus to direct Einstein's considerable powers of imagination to the problems posed by time and motion in late-19th century physics (Galison 2000).

modeled as a quest to find the highest peak in a randomly generated three-dimensional space populated by "hills" of varying height. By themselves, the knowledge-seekers in these simulations could only sense the relative elevation of a small space in their immediate environment – just like we as individual "knowers" only possess a small fraction of all the potential knowledge that might be possessed in any epistemic domain. This is what makes our own quest for knowledge, like that of the “agents” programmed into Zollman's simulation, a “bandit problem.” Also like us, Zollman's agents were not limited to their own epistemic resources; instead, they could become aware of the findings of others. It is this awareness of the findings of others Zollman manipulated to create different test conditions in his simulations. In some test conditions, the agents were quickly made aware of the findings of all the other agents; in others, they remained ignorant of most others and were left to explore on their own or on the basis of the knowledge accumulated by their own isolated "community" – i.e. other "climbers" near them on the exploratory plane – for a longer time. The crucial finding for our purposes is this: in Zollman's work, the simulations in which agents' access to the findings of other agents was *more limited consistently outperformed* the simulations in which agents are more aware of the findings of other agents. That is, the epistemic communities populated by comparatively *more epistemically isolated* agents consistently found higher peaks, and found them more quickly, than did those in which knowledge was shared more quickly and fluidly.²⁴⁵ Drawing on these computer simulations and his examination of the history of science through the example of PUD, Zollman concludes that epistemic communities faced with "bandit problems" may stand to benefit from at least some of their members behaving "in ways that make each individual look epistemically sub-optimal" because "when viewed as a community, their behavior becomes optimal."

This disparity between *individual*-level rationality and *collective* epistemic outcomes, dubbed the "Independence Thesis" by some²⁴⁶ and the "Zollman effect" by others,²⁴⁷ has been the subject to vigorous debate among philosophers of science and social scientists for the last decade, with some finding support for the thesis in the form of other computer simulations²⁴⁸ as

245 Ibid., 28-30.

246 Mayo-Wilson et al (2011)

247 Rosenstock et al (2017)

248 Douven (2019)

well as experiments involving actual human decision-making groups,²⁴⁹ while others have cited their own simulations and/or experiments to suggest that the Independence Thesis may only apply to a narrow range of circumstances.²⁵⁰ While this debate is important and fascinating in its own right, the important things to note for our purposes are that (a) no one contests Zollman's original argument that the Independence Thesis applies *within some range* of knowledge-seeking circumstances, (b) this range of applicable circumstances includes at least *some* human knowledge-seeking endeavors, and (c) we are by all indications very far from having anything like a precise theory of how to apply the principle to human affairs in the real world. My argument is not that the Independence Thesis *always* applies, nor that we can point with any appreciable degree of accuracy to some actual group decision-making scenario and state that the Independence Thesis applies to it. Instead, it is that because we know the Independence Thesis to apply to *some* range of actual human knowledge-seeking situations, we cannot straightforwardly infer that undesirable collective epistemic outcomes will result from what, applied to individuals, may be aptly described as epistemic vices, such as the dogmatism and ideological isolationism Progressives describe as typical of citizens whose online behavior takes place in “echo chambers” and “filter bubbles.” As Zollman puts it, “when analyzing particular behaviors of scientists (or any epistemic agents) we ought to think not just about the effect their behavior has on their individual reliability, but on the reliability of the *community as a whole*.”²⁵¹ The implication for the purposes of the current discussion is that even if “echo chambers” and “filter bubbles” are as common online as Progressives suggest and even if carrying out one's online existence primarily in them is properly described as an individual epistemic vice, the epistemic well-being of democracy as a whole might *still* be better off than if all citizens were routinely exposed to the full breadth of ideas, opinions and information available online.

To explain *why* this might be the case, I will now turn to Gerken's essay “Second-Order Diversity.” The central argument in Gerken's essay is that although legal scholars have been mainly preoccupied with what Gerken calls “first-order diversity,” their focus on *that* kind of diversity leads them to embrace public policies that may actually undermine the aims they see themselves as pursuing, and that we therefore need to think more broadly about what *kind(s)* of

249 Lazer et al (2018)

250 Mason and Watts (2012), Rosenstock et al (2017)

251 Ibid., 33, my emphasis

diversity best contribute(s) to those aims. For Gerken, the first-order diversity with which legal scholars have chiefly been preoccupied is nearly synonymous with *proportionality*. It aims for diversity *within* each decision-making group. Second-order diversity, on the other hand, does not worry too much about whether each *individual* group is reflective of the populace as a whole; in fact, it prefers for there to be a diversity *of* those groups in terms of their internal composition, so long as the *overall collection* of groups winds up more or less accurately reflecting the composition of the whole.

Gerken uses the examples of jury selection and the drawing of congressional districts to illustrate the meaning and implications of these concepts. Our current way of selecting jurors in the United States roughly approximates random assignment. Since this means that “a number of juries will look nothing like the population from which they are drawn,” i.e. not be a statistically proportional representation of the composition of the population as a whole, many policy makers and legal scholars have objected to the system unfair and proposed things like “racial quotas, stratified selection procedures, channeling litigant choices, and jury subdistricting” to make sure that *each jury* is sufficiently reflective of the populace as a whole.²⁵² In doing so, those policy makers and scholars are pursuing *first-order* diversity. While Gerken acknowledges that random selection does not foster first-order diversity in this way, she points out that *over time and on the whole* random selection *does* encourage proportional representation, even though *many individual juries* will not be proportionally representative. Because it ignores the proportionality (or not) of *each individual* group but manages to approximate proportionality *on the whole*, randomized jury selection is an example of *second-order* diversity. A similar distinction can be seen in debates about how to fairly draw the lines of congressional districts. A method emphasizing first-order diversity would draw each district such that every individual district contained a proportion of every relevant group equal to that of the relevant political unit being subdivided into districts (such as an individual state or the country as a whole). However, for a long time after the passage of the Voting Rights Act, legal scholars preferred the creation of “majority-minority” districts, which resulted in *some* districts being disproportionately White (for example), while others were disproportionately Black, Latino, etc. Legal scholars preferred majority-minority districts rather than drawing uniformly proportional districts because they

252Gerken (2005), 1115-6

viewed this as the best means of improving the chances of those minorities getting elected to political office and improving the diversity of *that* body as a whole. Since it ignores or even encourages the development of individual groups that are not proportionally representative of the population as a whole while at the same time aiming to foster diversity representative of the populace as a whole at a broader level, majority-minority districting is a means of drawing electoral districts that emphasizes *second-order* diversity.

The relevance of Gerken's argument to this chapter stems from her consideration of the ways the effects of policies emphasizing first- or second-order diversity can easily produce counter-intuitive *collective* effects. More specifically, Gerken points out that in both the example of jury selection and district drawing, insisting on diversity *within each individual group* can wind up undermining diversity *on the whole*. Consider jury selection in a district that is 60% White, 25% Latino and 15% Black. If every individual jury is a proportional representation of the populace, that means *every single jury* will be majority White. Usually, legal scholars worried about the implications of random selection are worried about the *under-representation* of minorities, but if juries were selected in a way that adhered to first-order diversity that would be a guarantee that *no* individual jury in such a district could ever possibly be composed of a majority of Blacks, Latinos, or even a combination of both. This could easily be problematic in precisely the cases legal scholars are most worried about – that is, cases in which racial prejudices common to the majority group are likely to affect their verdicts, as has no doubt often been the case in American history even up to the present. If such prejudices are rampant, first-order diversity would appear to be more likely to guarantee that *every* trial is influenced by those prejudices, whereas by allowing some subset of juries to be disproportionately Black or Latino would at least allow *some* trials – a number roughly proportional to the prevalence of Blacks and Latinos in that district – to be influenced most heavily by the judgment of those minority groups. A similar principle applies to the debate about how to draw electoral districts. If a state, for example, draws every single district within its boundaries in a way that results in each district being composed of the *same* proportion of each group that lives within its boundaries, the end result will be that whatever group comprises the majority of the state as a whole also comprises the majority of *each individual district*. In both examples, placing a strong emphasis on maintaining diversity *within each individual group* actually *undermines* diversity at the collective

level,²⁵³ thus potentially undermining the very aims sought by those advocating policies that emphasize first-order diversity.

I believe similar tradeoffs apply when we consider the relationship between the sorts of ideological, identity, or other discursive groups that form “echo chambers” and “filter bubbles” online and the epistemic well-being of democracy as a whole. The internet and social media in their current forms have enabled lay citizens all across individual countries and indeed around the world to come into contact with one another and, in some ways, *avoid* contact with those not like them with an ease not remotely approximated in prior eras. Undoubtedly, plenty of what motivates many of the members of these individual groups is the sharing of a common identity, ideology, or other set of preferences that motivates them to affirm one another's perspectives and preferences while disparaging those they identify as their opponents. This drives the *internal* composition of these individual groups to be non-representative of society as a whole and is therefore not conducive to first-order diversity. Though I have said that I believe this tendency to be exaggerated by Progressive commentators, I do not deny that it is a phenomenon that occurs regularly both online and offline. To the degree that “echo chambers” and “filter bubbles” *do* characterize an individual citizen's online environment, they are screened from alternative points of view in ways that serve to protect their current worldview, limiting their chances of improve their ideas. That is near to the very definition of bad in individual epistemic terms. Progressives like Cass Sunstein argue that it is also “undesirable from the democratic standpoint.”²⁵⁴ In Sunstein's view, this is because the “greater ease and frequency” with which “like-minded people are connecting...with one another” without “hearing contrary views” is likely to lead to a “high degree of fragmentation, as diverse people...end up in extremely different places simply because of what they are reading and viewing,” as well as a “high degree of error and confusion” because the sort of content that proliferates in online platforms like Youtube “will lead like-minded people to end up with a distorted understanding” of reality.²⁵⁵ Sunstein also reflects the common Progressive preference for a media environment dominated by legacy media companies, which he calls “general-interest intermediaries,” whose need to appeal to a wide audience meant that they would “expose people to a range of topics and views” not particular to the identity or

253Gerken (2005), 1124-1126

254Sunstein (2016), 66

255Ibid., 76-77

ideological group.²⁵⁶ Sunstein's preoccupation with the dangers that might arise from allowing so many ideological communities that are *internally* non-representative of the general populace to proliferate in the online environment and his preference for a legacy news media-dominated informational environment in which a *single kind* of diversity – diversity between the relevant viewpoints, as assessed by professional journalists – are all indicative of a *first-order* perspective on the type of diversity that is beneficial to and indicative of the epistemic well-being of democracy. This perspective is analogous to the perspective held by those who advocate for quota-driven jury selection and proportionally drawn electoral districts and is, accordingly, subject to the same danger of *undermining* diversity at the collective level precisely *through* the pursuit of first-order diversity at the level of each individual group. Indeed, I argue this was precisely the case during the broadcast era treated with such veneration by Progressives. The informational environment dominated by the legacy news media between the 1950s and 1990s was one in which journalists claimed, many of them sincerely, to follow a rigid code of ethics designed in part to guarantee that the legacy news was inflected with some degree of perspectival diversity, for example by newspapers including quotes taken from “both sides” in political stories and television stations giving equal air time to politicians from the two major parties. As Sunstein himself admits, however, subsequent empirical research strongly indicates that this practice led news producers to “avoid controversial issues entirely and to present views in a way that suggested a bland uniformity.”²⁵⁷ Not only that but, as epistemologist David Coady points out, the type of diversity the legacy news sought to incorporate into its reporting often reduced to implying that there were “precisely two sides to each political issue which should each be given equal time” because “political power in most Western countries...takes the form of a duopoly.”²⁵⁸ As a result, points of view that were not endorsed by one or both of these two powerful groups were frequently ignored or, if not that, often portrayed by the legacy news as “extremist” positions – a term which, it is worth noting, Sunstein and other Progressives also frequently employ. Is it a coincidence that the era of greatest dominance of the informational environment by a legacy news media following such practices runs hand-in-hand with the rampant adoption of anti-communist, anti-socialist and pro-capitalist sentiment – all views subscribed to with near

256Ibid., 43

257Ibid., 84

258Coady (2011), 148

universality by Republican and Democratic political leaders throughout the Cold War era – in the United States public, or that it was only *after* 21st century communications technologies allowed the fostering of new voices into citizens' informational spheres that a self-described socialist, Bernie Sanders, was able to muster widespread electoral support – support that relied heavily on the internet and social media while Sanders was largely ignored in favor of mainstream rival Democrat Hilary Clinton? I do not think it is. Instead, I think it is the likely result of relying on a first-order model of diversity. By emphasizing that *each knowledge-seeking or decision-making group* (each TV station, newspaper, etc.) seek to achieve *the same kind* of diversity, the broadcast era of media stifled diversity at the *collective* level in comparison with the internet era in which we now live. By assuming “echo chambers” and “filter bubbles” are bad for democracy because they diminish the *internal* perspectival diversity *within* each discursive group online and defending an informational environment dominated by “general-interest intermediaries,” Sunstein and Progressives risk embracing policies that stifle collective diversity in the same way, such, for example, as suggesting that Google and Facebook should tweak users' algorithms so that users are exposed to (what Facebook and Google determine to be) “real news” or “the truth” or “non-extremist views.” The upshot of these considerations about first- and second-order diversity is similar to the upshot of Zollman's examination of the relationship between the conditions of individual versus collective knowledge-seeking. In both cases, policies that seem intuitively correct when viewed at the level of individuals or groups within a larger collective sometimes prove detrimental to good outcomes at the level of collective as a whole. Likewise, in both cases, when we apply this logic to the topic of “echo chambers” and “filter bubbles,” we discover reason to hesitate before embracing Progressives' argument that the (supposed) rise of these things is detrimental to democracy.

Indeed, I believe we have good reason to suspect the truth may be just the opposite – that lay citizens' increased ability to find people like them online and form mutually supportive communities through digital technologies is a net *benefit* to democracy. Sexually “deviant” communities such as the gay, trans and “kink” communities are plausible examples of just this phenomenon. In the pre-internet era, people who had these sexual orientations had little to no access to like-minded others and were often forced to live “closeted” lives for fear of social rejection and even physical abuse. When the internet arose, it offered to these and numerous

other wrongly discriminated against communities the opportunity to seek out and receive support from people who shared their lifestyle, views and preferences. Their ability to police their own communities and create “echo chambers” to keep out other voices and amplify their own helped them gather momentum by sharing ideas and affirmation in an environment in which their voices could not be as easily drowned out as they were in mainstream society. This allowed the individuals within those communities to gain confidence, to realize they were not alone, and to begin to speak with an increasingly unified, and increasingly powerful, voice. This increasingly powerful voice has, in turn, led to a steady increase in societal acceptance of these lifestyles which is more and more being incorporated into both laws and the policies of individual businesses, such as the increasing use of unisex bathrooms and an ever-diminishing tolerance for employers who discriminate against people subscribing to these lifestyles. Insofar as the process of online community-building and boundary-policing I described above has helped usher in these changes, 21st century communications technologies and their power to allow individuals to self-select their discursive communities and form echo chambers and filter bubbles has contributed to the well-being of our democracy, both epistemically and generally. It has helped greater society realize just how many “deviant” individuals there are and encouraged many of us to expand our views about what is “normal” and what is permissible.

I do not think this is a unique example of the benefits global democracies have already reaped as a partial result of internet-era citizens' ability to form like-minded communities of mutual support with the ability to exclude other points of view, i.e. “echo chambers” and “filter bubbles,” and will offer further examples in future chapters. For now, I will content myself with the argument as it stands. While the ability to self-select one's discursive community online and as a result form “echo chambers” and “filter bubbles” may in many instances, including for many of the sexually “deviant” citizens mentioned above, be indicative of or may itself be a manifestation of individual epistemic vice, that does not mean an increase in such ability equates to an increase in *collective* epistemic vice for democracy. On the contrary, in many important cases it is just the opposite.

Section Five: Erring on the Side of Doxastic Hubris

Individual Epistemic Vice and the History of Liberal Democracy

The importance of these analyses of the three conversations mentioned above extends beyond the potential for these conversations' pessimism to contribute to premature and/or mistaken diagnoses about the epistemic well-being of democracies in the early 21st century. The complicated and often counter-intuitive relationship between individual epistemic vice and collective epistemic virtue also bears profound implications for politics historically, theoretically, and practically. Historically, recognizing the ways individual epistemic vices can be collective epistemic virtues offers a way of explaining the historical viability of liberal democracy as a political system, a phenomenon which has befuddled a number of political scientists over the course of the past century as they have become more familiar with the low levels of political knowledge and high levels of individual epistemic vice of the sorts mentioned above that appear to be quite common among lay citizens across eras. Political thinkers have generally responded to this odd juxtaposition of rampant individual epistemic vice among liberal democratic citizens and the historical success of the liberal democracies composed of such citizens either by positing a set of compensatory individual epistemic habits (heuristics) that paint individual citizens as more virtuous epistemic actors than they may at first seem²⁵⁹ or by emphasizing non-epistemic factors in their explanations of liberal democracy's historical success²⁶⁰ or, more commonly, some combination of both.²⁶¹ While a wealth of insight can be gleaned from the ruminations of the scholars who have offered these responses and I believe there to be much merit in many of their arguments, I also believe that the thoughts offered in this chapter can add an additional, and needed, supplement to them. That is, *in addition to* citizens' ability to use efficient cognitive tools like heuristics to navigate complex informational environments *and* liberal democracy's success stemming in large part from features that have nothing to do with the epistemic habits of citizens, the historical success of liberal democracy stems partly from the ways it allows the individual epistemic vices of lay citizens to become collective epistemic virtues. By allowing individual citizens animated by motivated reasoning to defend their own pet theories, obstinately held

259E.g. Popkin (1991)

260E.g. Lippmann (1925), Przeworski (1999)

261Examples abound. See Schumpeter (1975 [1942]), Downs (1957), Schattschneider (1960), Posner (2003)

dogmas and irrationally embraced conspiracy theories in public, liberal freedoms of conscience and expression counteract our collective tendency toward complacency and conformity and ensure that dominant ideas are being questioned and venerated authorities are being kept watch over. And by giving each individual citizen a say in the selection of leaders and/or public policies, democracy offers a standing means of putting whatever contestation may arise from those ideas gaining currency among large sectors of the public into concrete political action. Part of liberal democracy's success as a political system, I believe, stems *from the very epistemic vices of individual citizens themselves*.

Why it is Better to Err on the Side of Doxastic Hubris than Doxastic Submissiveness

While improving our understanding of the political past and interpretations of the political present are, of course, important for a host of reasons, the main implications I want to emphasize from the arguments presented in this chapter have to do with our models of good citizenship and the public policies we see as appropriate for managing (or not manage) 21st century communications technologies. In short, I believe the points made in this chapter about the ways individual epistemic vices can be collective epistemic virtues strongly support doxastic self-rule as a collective epistemic virtue, and that it will be made best use of in liberal democracies as defined in Chapter One.

When one examines the examples of individual epistemic vice that can serve as collective epistemic virtues given in this chapter, a common thread can be seen running through all of them. This is that the individual epistemic vices-turned collective epistemic virtues entail situations where individuals rely on their *own judgment* in situations where that judgment is *flawed*. Motivated reasoning, for example, is a case where an individual's judgment is flawed by being biased in favor of some conclusion they *prefer* to be true, as may plausibly have been the case for either/both of Priestley and Wegener and is undoubtedly true for many of us much of the time. Likewise, anti-vaccers' opposition to vaccines is a case of their relying on their own flawed judgment, either about the facts of the matter regarding vaccine theory or at least in their choice of whose testimony about the risks and rewards of vaccines to trust. The phenomena of “echo chambers,” “filter bubbles” and dogmatism can be interpreted in the same way: to the extent that isolation from other views is the result of the choices of the individual agents – for example,

through their choices of which news to read, which groups to join, which people to “follow,” and who to “unfriend” online – it may be interpreted as those agents choosing to rely on their own judgment by expedient of depriving themselves of access to the judgment of others; similarly, dogmatism is excessive stubbornness in clinging to one's own belief even if one *has* allowed oneself to be exposed to evidence to the contrary. In all these examples of individual epistemic vice, individual knowledge seekers insist on relying on their *own judgment*, rather than deferring – or even *allowing themselves to be exposed to* – alternative arguments and points of view.

Another way of putting the point is to say that all these cases of individual epistemic vice imply a certain *excess* of *doxastic self-confidence*. I prefer this way of putting it because it aligns well with the language of virtue, which I have used throughout this chapter and which I believe is an instructive way of approaching civic epistemology generally. One of the hallmark characteristics of thinking in terms of virtues is to acknowledge that determining whether a particular human behavior is praiseworthy or contemptible depends not only on the behavior's *kind* but also its *degree*. Aristotle famously captured this principle when he wrote that virtue “is the mean by reference to two vices: the one of excess and the other of deficiency.”²⁶² This applies just as well in thinking about the role of what I have called doxastic self-confidence, or the insistence on maintaining one's own beliefs rather than deferring to the claims, arguments or evidence presented by others who disagree with us, as it does to the sorts of ethics discussed by Aristotle millennia ago. Doxastic self-confidence is an integral part of sound belief formation at the individual level. A *deficiency* of doxastic self-confidence – what might be called *doxastic submissiveness* – invites citizens to be easily duped by swindlers, charlatans, abusive relationship partners, used car salesmen, and demagogues, and thus to adopt unreliable beliefs and choose unwise actions based upon them. Yet, as we have seen in this chapter, an *excess* of doxastic self-confidence – what might be called *doxastic hubris* – is also a vice, causing them to cling too tightly to poorly supported beliefs and resulting, again, in unwise life choices. Thus, at the individual level, doxastic self-confidence can be both a virtue and a vice. Whether it is one or the other is a matter of degree.

What is the *appropriate* level of doxastic-self confidence for democratic citizens to manifest in dealing forming their beliefs in the complicated epistemic environment of the 21st

262 *Nichomachean Ethics* 1107a

century? Alas, answering this question precisely is no more possible today than it was for Aristotle in ancient Greece. On the contrary, Aristotle was correct when he insisted that since individuals differ in both personality and judgment and since the situations they face are so diverse, it is impossible for the theorist to formulate a universally applicable rubric for virtuous behavior. That is not to say, however, that we are entirely helpless when it comes to offering practical guidance to those who wish to be virtuous. On the contrary, while Aristotle's argument that virtue lies in-between two extremes is the most famous part of his *Nicomachean Ethics*, I believe the most practically useful line of argument in that book can be found immediately afterward. It begins by noting that virtue often lies closer to one of a behavior's vicious extremes than the other²⁶³ and that erring on the side of one of the extremes is often better than erring on the other. The upshot of all this, he writes, is that the "first concern of a man who aims at the median should, therefore, be to avoid the extreme which is more opposed to it." In other words, "since it is extremely difficult to hit the mean, we must, as the saying has it, sail in the second best way and take the lesser evil."²⁶⁴

I will not in this chapter examine the implications of Aristotle's arguments in determining whether it is appropriate at the level of *individual* epistemology to err on the side of doxastic hubris or submissiveness, though the next two chapters will imply the former. Here, however, I just want to argue that at the level of *collective* epistemology it is better for citizens to err on the side of doxastic hubris than doxastic submissiveness. My general reasoning for taking this position has already been suggested in Chapters One and Two and will be a central theme in the rest of this dissertation. The main reason we should encourage citizens to err on the side of doxastic hubris rather than doxastic submissiveness is that the assertive exercise of doxastic self-rule by lay citizens is essential to our hopes of *improving our collective knowledge* and *uncovering and resisting illegitimate authority*, tasks which I take to be of paramount social and political importance for polities in all places and at all times. One implication of this chapter is that citizens exercising doxastic self-rule *too* assertively (doxastic hubris) can still be actively contributing to the epistemic well-being of democracy. I do not believe a similar relation is true

263 "In some cases it is the deficiency and in others the excess that is more opposed to the median" (Ibid., 1109a). As the translator in this case has correctly inferred, this means that Aristotelian virtue is more correctly conceived as a "golden *median*" than a "golden *mean*."

264 Ibid., 1109a-1109b

when citizens err on the side of doxastic submissiveness. On the contrary, I believe that sort of excess hampers our collective ability to improve our beliefs and actively invites the development of primary bads like corruption, whose dangers far outweigh potential benefits like increased social cohesion.

Implications for Public Policy and Civic Epistemology: Free Expression and Doxastic Self-Rule

As in the last chapter, in this chapter I have declined to object to certain dominant themes in contemporary conversations about the epistemic impact of 21st century communications technologies on democracy. I have not pressed the argument that the “echo chambers” and “filter bubbles” worried about by Progressives seem to be less common than Progressives tend to imply, nor that the epistemic behaviors labeled as “cognitive biases” are less widespread, less intractable, and less deserving of the label of epistemic vice than Progressives treat them as being. Though I believe there is some merit to all of these objections and (what I take to be) the mistakes of Progressives on these matters are important, I have not chosen to emphasize these objections. Instead, I have so far restricted myself to critiquing the ways Progressives tend to *think about what it means* for democracy to be doing well epistemically and the sorts of dispositions and behaviors on the parts of lay citizens that are likely to contribute to democracy's epistemic well-being. I think these matters are of great practical importance in two ways. First, they bear implications for public policy, especially with regard to the ways we choose to require or encourage authorities to, or discourage or forbid authorities from, managing the unprecedentedly powerful, cheap and widely accessible channels of communication brought forth by the technical advances of the 21st century. Second, they affect the behaviors and dispositions we encourage (or discourage) in the citizenry as part of their civic duty. I believe the arguments presented in this chapter favor a less regulated communications environment and a more doxastically assertive lay citizenry and will close this chapter by briefly explaining why.

In this and the preceding chapter I have argued *both* that (a) Progressives' deep pessimism about lay citizens is at best premature and worst deeply misguided and (b) even many of the same behaviors from which Progressives correctly draw pessimistic conclusions about lay citizens at the level of *individual* epistemology can serve as a basis for optimism about the epistemic well-being of democracy at the level of the *collective*. Both of these points undermine

the increasingly common Progressive position that democracies' informational environment needs to be regulated by some set of authorities, be it a government-appointed panel or some set of “fact-checkers” employed by the likes of Facebook and Google, in order to prevent a collective epistemic collapse. To the degree that lay citizens' judgment is better than Progressives depict it as being, there would seem to be less need for its informational environment to be managed for them by some set of authorities. In addition, even in cases where their epistemic behavior *is* sub-optimal at the individual level, including at least a few of the specific instances most consistently cited by Progressives as reasons why the informational environment needs to be managed, that behavior could *still* be beneficial to democracy in the long run. To the degree that lay citizens' information is curated by authorities, democracy is deprived of *both* the standing benefits of doxastic self-rule for collective knowledge growth and the uncovering of and resistance to illegitimate authority *and* the counter-intuitive collective benefits of individual citizens' epistemic vices. The benefits of doxastic self-rule are proportional to the degree to which individuals' judgment determines their beliefs. Since, as Mill pointed out, the act of an authority deciding what information citizens are *allowed* to be exposed to “exclude[s] every other person from the means of judging,”²⁶⁵ it necessarily diminishes the potential for society to reap the benefits of individuals' exercise of doxastic self-rule. The arguments presented so far, therefore, militate against the sort of regulation of the informational environment increasingly advocated by Progressives.

If part of the danger of such increased regulation stems from its neutering of the epistemically beneficial role of doxastic self-rule by lay citizens, then those benefits will also be diminished to the degree that we discourage lay citizens from *exercising* doxastic self-rule. My objection to Progressives' approach to the epistemic issues brought on by 21st century changes in communications technologies is not limited to opposing their increasing inclination toward top-down management of the informational environment. It extends to the behaviors and dispositions they encourage (and discourage) in citizens who engage with that informational environment. Progressives' fixation on negative examples of citizens' cognitive habits and behavior online leads them to depict citizens as fundamentally untrustworthy epistemic agents, at least when released into the “wilds” of the hitherto largely unregulated informational environment of the

²⁶⁵Mill (2003 [1859]), 78

internet. This depiction of citizens as fundamentally untrustworthy epistemic agents stands in stark contrast to Progressives' depictions of their favored epistemic authorities – prominently, the scientific community, the legacy news media and the ill-defined set of self-proclaimed epistemic authorities known as “fact-checkers” – all of which they argue or strongly imply to be fundamentally trustworthy, to the point of being guardians²⁶⁶ of democracy. The overall image is one in which 21st century democracy is imperiled by lay citizens *precisely because* lay citizens have unprecedented *access to* information, as well as the unprecedented ability to spread *their own* ideas, as opposed to that access being restricted and that ability being reserved to a select few.²⁶⁷ In other words, democracy is *threatened by* the judgment of the lay public and the solution involves *deferring to authority*. In addition to its tendency to lead Progressives to embrace top-down management of the informational environment as a matter of public policy, this image leads Progressives to embrace a model of civic epistemology based on doxastic submissiveness. Were citizens to adopt the Progressive view, they would see *themselves* and *each other* as epistemic agents whose judgmental capacity is so compromised as to render them not just largely unfit to form reliable opinions about what is true on a wide range of politically important topics, but indeed so unfit that they would be *imperiling democracy* simply by attempting to do so. It is a short step from such a view to the conclusion, also strongly implicit in the Progressive position, that responsible citizenship consists largely in *resisting* the temptation to assertively exercise one's own judgment, and of making a habit, instead, of routinely deferring to authorities that (one supposes) “know better.”

I believe the adoption of such a disposition by large subsets of the population would be damaging to democracy in the same way as the widespread imposition of “fact checkers” to manage citizens' informational environment from above, and for the same reasons: first, society depends on citizens' assertive exercise of doxastic self-rule for the growth of knowledge and the discovery of and resistance to illegitimate authority; second, there is much reason to suspect the view of citizens on which Progressives' recommendation of doxastic submissiveness is built is at excessively pessimistic; third, even in many cases where a pessimistic interpretation of

²⁶⁶The double entendre here is intentional.

²⁶⁷“Freedom of the press is guaranteed only to those who own one,” wrote Abbott Joseph Liebling in 1947

(Liebling 1947). Progressives very often seem to wish to return to an era, like Liebling's, where ownership of a press – i.e. a means of getting one's thoughts published and disseminated widely – was restricted to a very select few.

individual epistemic behavior is justified, that very behavior can be beneficial to the epistemic well-being of the collective. It is better, therefore, to emphasize the importance of citizens being assertive in their exercise of doxastic self-rule, even if that means inviting the vice of doxastic hubris, than it is to risk their erring on the side of doxastic submissiveness by over-emphasizing their cognitive unreliability and the comparative superiority of epistemic authorities. One reason is that even when such emphasis leads citizens to the vice of excess at the individual level, that vice may well to be a virtue when it comes to the collective.

Conclusion: The Need for a New Model of Epistemic Well-Being

In these first three chapters, I have attempted to identify and object to the *deep pessimism* about lay citizens implicit (and often enough explicit) in Progressives' interpretation of the epistemic impact of 21st century communications technologies on democracy, as well as the *deferential* and/or *submissive* model of civic epistemology built largely on that pessimism. In Chapter One, I emphasized the importance to liberal democracy of citizens' willingness to exercise *doxastic self-rule* and its implied *cautious optimism* about lay citizens' judgment, and warned that Progressives' lack of support for either position is likely to lead to their abandonment of the liberal component of liberal democracy. In Chapter Two, I raised a number of objections to the ways commentators have tended to view the impact of 21st century communications technologies on democracy's epistemic well-being, ranging from empirical objections to the ease with which commentators blame those technologies for undesirable features of our collective epistemic landscape, such as climate denial and the flat earth movement, to theoretical objections to commentators' excessively deferential and naive view of the proper approach citizens ought to take when assessing the claims of scientists, journalists and other would-be epistemic authorities vis-a-vis their opponents'. The way commentators tend to view the impact of 21st century technologies on democracy, I argued, tends to unfairly prejudice them against lowered barriers to publication, too easily settle for pessimistic views of lay citizens' activities online based on impressions rather than empirical evidence, encourage citizens to adopt an excessively deferential attitude toward conventional epistemic "authorities," and play down the ways 21st century communications technologies are helping *improve* the epistemic well-being of

democracy while over-emphasizing the threats those technologies purportedly pose to it. These tendencies are partly *caused by* their *prima facie* disposition of pessimism about lay citizens' capacities of judgment and are *cited as grounds for* further pessimistic conclusions about both citizens themselves and their fitness for liberal democracy, to both of which I object. Finally, in this chapter, I objected to commentators' tendency to assume that what they identify as epistemic vices common to individuals and groups in 21st century democracies are detrimental to the epistemic well-being of democracy as a whole, arguing that there are reasons both historical and theoretical to doubt this claim. This tendency has, once again, tended to predispose commentators toward pessimistic interpretations of the changes wrought on our society by the invention of new communications technologies, and is reflective of and contributes to the view that democracy succeeds *despite* its citizens, rather than because of them.

I believe there are common threads that run between these Progressive tendencies. By tracing these threads one to another, I believe we can begin to see patterns which, once woven together, constitute a conventional – and largely coherent – over-arching way of thinking about the epistemic well-being of democracy, and how the behaviors and characteristics of citizens contribute to, or detract from, it. As internally coherent as this view may be, I believe it is deeply flawed. The most serious flaws are that (a) the ordinary epistemic behaviors and capacities of lay citizens are excessively condemned, both in terms of their individual rationality and especially in terms of their contributions to collective epistemic well-being; and (b) the implications of chronic epistemic fallibility – the most important tenet of scientific and, I believe, any serious inquiry – are consistently under-appreciated. The result is for Progressives to consistently devalue, and increasingly attempt to actively restrict, the very sort of dissent integral to both the growth of collective knowledge and our ability to discover and resist illegitimate authority. I have argued that we have good reason to resist this way of thinking, both because it is founded on an excessively pessimistic view of the judgment of lay citizens and because, partly because of that pessimism, it advocates an attitude of doxastic submissiveness that is *itself* likely to undermine the epistemic well-being of democracy. At least so far, I have argued, we have no reason to suppose the cautious optimism implicit in support for liberal democracy is ill-founded, and accordingly we ought to consider epistemic self-confidence as a *civic virtue*, and encourage citizens to err on the side of exercising more, rather than less, doxastic self-confidence.

In each of these first three chapters, I have started with a broad view of the phenomenon investigated – liberal democracy in Chapter One, the epistemic well-being of democracy in Chapter Two, and collective epistemic virtue in Chapter Three – and drawn inferences about the sorts of behaviors and dispositions we ought to encourage in citizens from those macro-level considerations. In the next two chapters, I will take the opposite approach, beginning by looking closely at belief formation under 21st century conditions from the perspective of the individual citizen and extrapolating broader conclusions from that starting point. I hope to demonstrate that the cautious optimism about the judgment of lay citizens and the model of civic epistemology based on doxastic self-rule I have advocated so far are just as well supported when we begin by looking at things from perspective of the individual citizen as when we take the opposite tack. That is, adopting a general attitude of epistemic self-confidence when it comes to deciding what – or, as we shall see, much more commonly *whom* – to believe is as justified at the level of *individual* as *collective* epistemology, even if, as we have seen, it can at times rightly be named a vice.

Chapter Four: Civic Epistemology

In Chapter Two I noted that the epistemic well-being of democracy in the 21st century is a central concern for Progressives. The most subscribed-to and frequently-reiterated credo that unites Progressives on this front is one with which all of us are familiar and which was repeated a few years ago in an interview by the historian Howard Zinn (among, no doubt, innumerable others). “Democracy depends on citizens being informed.”²⁶⁸ For America, the roots of this sentiment trace back to the days before independence. In a letter entreating the inhabitants of Quebec to join them in resisting the British Empire, the Continental Congress included among their arguments the need for a free press, whose functions included “the advancement of truth, science, morality, and arts.”²⁶⁹ Though the founding fathers were far from having a naively romantic view of the wisdom of the common man – one of the many reasons most of them scorned the idea of democracy – they *did* have an Enlightenment optimism that, given enough time and information, the public would get things right. For them, then, the freedom of the press was the best safeguard for the ability of the public to self-correct.²⁷⁰

Notably, the founding fathers did not live in the age of social media. For many, including many Progressives, the Information Age is calling into question the founding generation's faith in the curative effects of information itself. Such, at least, is the implication of much of the talk that has surrounded that prodigal child of the Internet, “fake news.” Kathleen Hall Jamieson warns of the “Demise of 'Fact' in our Political Discourse.”²⁷¹ The *Tampa Bay Times* calls fake news a “scourge on informed citizenship.”²⁷² The editorial board of the *New York Times* has called on

268It is worth noting that while Zinn can accurately be described as “Progressive” in the sense that the term “Progressive” is attached to those sympathetic with the plight of the working class, in favor of redistributive governmental policies, etc. However, in the interview containing this quote Zinn deviates from one of the central planks of Progressivism as I have defined it, as the quote is lifted from a passage in which Zinn criticizes the legacy news media for failing to keep the people informed. Indeed, Zinn accuses the legacy news media as performing so poorly that “we cannot really say we have a functioning democracy” (Sirucek 2010).

269“Continental Congress to the Inhabitants of Quebec” (1774)

270Among the founding fathers, Jefferson wrote the most on the topic of freedom of the press. His position of simultaneous scorn for the papers of his time and staunch defense of the value to society of leaving them free to publish their rubbish is something I shall cover in some depth further on in this dissertation.

271Jamieson (2015)

272*Tampa Bay Times* (2016)

Facebook to take an active role in vetting and blocking fake news, arguing that it “owes its users, and democracy itself, far more” than what it has so far done to vet fake news.²⁷³ Public figures have chimed in on the topic, with the apparent consensus being that voiced by Apple CEO Tim Cook, that fake news is “one of today's chief problems,” necessitating that Apple and other information-driven companies take steps to curb its spread.²⁷⁴ After heavy public pressure and much initial resistance following the 2016 presidential election Facebook, Twitter and Google have all pledged to take steps to combat fake news²⁷⁵ and have taken increasingly proactive stances toward combatting it through a variety of measures, some designed to inspire readers to pause and engage in critical thinking and some to prevent the content from ever reaching users' eyes in the first place. Likewise, governments around the world have stepped up measures to combat fake news, with France having already passed²⁷⁶ – and used²⁷⁷ – legislation against the dissemination of (what the government deems to be) “fake news” in the lead-up to elections and similar legislative moves being enacted in numerous other countries. A growing consensus, at least among public-minded scholars and commentators in the public sphere, seems to be emerging that fake news and misinformation are not just problems but a scourge, an unprecedented assault on the arteries of information on which our civic body so vitally depends, and that some set of authorities needs to step in and manage the informational environment lest it corrode the very pillars of democracy.²⁷⁸

In the last two chapters I have critiqued the incomplete and in some ways misguided view of what it means for democracy to be “doing well” epistemically on which these worries, and the notions of civic epistemology derived from it, are based. Against the Progressive view, in Chapters One through Three I have defended a certain *modest optimism* about lay citizens' judgment and the encouragement of *doxastic self-rule* by citizens. The primary normative concerns in these chapters have been expressed in terms of the effects of citizens' practice of doxastic self-rule on *collective* outcomes, primarily the growth of our *collective* knowledge and our *collective* ability to discover and resist illegitimate authority. In this chapter and the next, I

²⁷³New York Times (2016)

²⁷⁴Good Morning Britain (2017)

²⁷⁵Wingfield, Isaac and Benner (2016)

²⁷⁶Library of Congress (2020a)

²⁷⁷Jamieson (2018, 11-12)

²⁷⁸For a much more extensive library of citations demonstrating the impressive extent to which this series of positions holds sway among intellectuals and in the public sphere, see Farkas and Schou (2020), ch. 3

will home in on the individual by taking an in-depth look at the epistemic circumstances faced by *all* of us who, as individuals, are trying to figure out what to believe in the large, complicated, interconnected world of 21st century politics. In this chapter, I will argue that any convincing model of normative *individual* epistemology must be based on *doxastic self-rule*, or the demand that one's beliefs be based on their having passed a test of scrutiny in the lights of one's own judgment. Thus, starting out from the standpoint of individual normative epistemology leads us to precisely the same conclusion as I reached in the first three chapters, when I started out from the standpoint of collective normative epistemology: the exercise of doxastic self-rule by individuals must be a core component of *any* persuasive account of normative epistemology, no matter whether we begin from the standpoint of the individual or the collective.

The argument will proceed as follows. In Section One, I will introduce core concepts to be used in this and the following chapter and define important terms. In Section Two, I will argue that the active exercise of judgment is the essential determinant of whether or not an individual's beliefs are *justified*, and therefore must play a central role in any normative account of individual belief formation. In Section Three, I will focus on one particular way we go about acquiring knowledge, through *trust in testimony*, which, I will argue, is the proper point of focus for contemporary civic epistemology since it is the source of the overwhelming majority of our beliefs and of special concern in the Information Age. Because the world is large and complex, for most of our beliefs, including many important beliefs about politics, we do not have the relevant experience to evaluate a given proposition on its merits. The fact that lay citizens' judgment is often insufficient to allow them to reliably evaluate propositions on their own merits does *not* mean individuals' judgment is irrelevant or inapplicable in the decision about what to believe in such cases; it merely means that the proper role of judgment in such cases is to evaluate *which testifiers are legitimate epistemic authorities* on the matter at hand. In Section Four, I will argue in favor of a civic epistemology based on *epistemic vigilance*, or the active exercise of judgment on the part of individual citizens, and will contrast this model of civic epistemology against the *selective skepticism* favored by Progressives. I will conclude the discussion by illustrating the advantages of adopting an approach based on epistemic vigilance over one based on selective skepticism with regard to the issue of fake news.

Section One: Definitions and Background

“News” vs. “The News”

One of the central concepts around which this chapter, as well as the next, will be organized is the concept of *news*. The terms “fake news” and its implicit opposite “real news” are invoked with great regularity in contemporary discussions about what is wrong with 21st century communications technologies and the post-truth movement spawned by them. Progressives usually seem confident that they, or at least authorities they are confident can be straightforwardly identified and trusted, are better positioned epistemically to determine which reports about the world are reliable than most lay citizens, whose judgment they mistrust. Very often, these reliable and trustworthy authorities include the “real news,” or legacy news media, while other forms of testimony, such as word-of-mouth, rumor, social media, and a multitude of online sources such as the “blogosphere,” are treated with suspicion and disdain. Because of this, Progressives advocate for citizens to adopt an asymmetric epistemic disposition to the claims of the “real news” vis-a-vis the claims citizens encounter via these other sources. Specifically, they encourage citizens to adopt a default attitude of trust in the claims of the legacy news media and a default attitude of skepticism toward the others. As we shall see, they also often embrace public policies that would give those they supposed to be “experts” about which news ought to be trusted the power to manage the informational environment through censorship, such as “fact-checkers” in the employ of mega-corporations like Facebook and Twitter.

As should be obvious by now, I could scarcely disagree more with this position. For now, though, I just want to clarify the various ways I am going to use the term “news.” As used in common language, “news” is different than “*the news*.” *The news* refers to a peculiar assortment of informational products gathered and disseminated by members of the profession of *legacy journalism*, which began to be developed in the 19th century yet did not really resemble the professional, code-governed enterprise we associate with that term until the mid-20th century. *News* without the “the” is a much broader, and much *older*, term. It encompasses *the news*, but also includes a tremendous range of other human testimonies. As news historian Michael Schudson points out, it is news, not *the news* that is in demand, for instance, in Shakespeare's

Richard III,²⁷⁹ where in different parts Gloucester demands it of Hastings, two gossiping citizens of London demand it of each other, the Duchess and Archbishop of York demand it of a messenger, and Buckingham seeks it out in the rarefied air of the Guildhall. The *news* these characters seek is *information of interest regarding recent but distant events*. In order to get that information, the characters make use of their *social resources* – Gloucester his well-traveled friend, one citizen her gossiping co-citizen, the Duchess a messenger hired to transmit testimony among the upper echelons of society, and Buckingham his fellow guildmates. *News* in this sense is still as important to us today as it was in the time of Shakespeare, and often has the following qualities: (a) the channels through which it disseminates are almost entirely *ad-hoc*, (b) the *means* by which it travels are *informal* and can be summed up colloquially (with a few charitable amendments to apply to the digital age) as *word-of-mouth*, (c) the proximal source is usually some *familiar acquaintance*, (d) their newsworthiness is determined only by their vague designation of *being of interest in the opinion of the discussants*, and (e) it is *believed by the discussants to be accurate*, if not in detail, at least in essence.

Of these features of the *news* sought by the characters in *Richard III*, the only ones that are consistently true of *the news* are (d) and (e), that is, the information is thought to be accurate and of interest to the discussants. In all other respects, it is quite different. In order to count as *the news*, information must be disseminated via through a *formal process* that is conducted by a *particular type* of source – a professionalized institution partly or wholly dedicated to discovering, producing and spreading its own special informational product. The *source* through which it disseminates is not generally some familiar acquaintance but rather a *distant institution* known only through reputation, though the reports of such an institution may themselves be shared through casual conversations between friends or strangers or, in contemporary times, through Facebook feeds or posts on Reddit or Twitter. Generally speaking, the difference between *news* and *the news* can be summed up by pointing out that *the news* is an *informational product originally composed and distributed by professional organizations*, while *news* need not be. Thus, all of *the news* is *news*, but not all *news* is *the news*.

²⁷⁹Schudson (1990), p. 3. Schudson does not use the precise terminology I adopt here, but his position is the same.

News: Any testimony made by one party to another party or other parties on *recent* events or phenomena perceived to be *of interest* to the communicating parties.

The News: An informational product revolving around recent events or phenomena which is gathered and disseminated by a set of professional organizations that arose in the 19th and 20th centuries and came to specialize in that process of gathering and dissemination.

“Real” vs. “Fake” News

Political scientist Jonathan Ladd has demonstrated that even though *news* without the “the” can have the broader meaning described in the preceding sub-section, usually when Americans hear the term “news,” their default assumption is that it means *the* news.²⁸⁰ I expect the same is true of English speakers in other parts of the world as well. Ladd's conclusions are further supported by the way the term “fake news” is often – though not always, as we shall see in chapter two – used to describe an informational product that imitates *the* news. The term “real news” is used much more seldomly than the term “fake news,” but it is easy to extrapolate what the former means just by taking the opposite of the latter. A “fake” of something is an imitation meant to pass as the “real” thing. Since “fake news” is an imitation of *the* news meant to pass as such, the “real” news must be what I have called *the* news. This gives us the definitions of “real news” and “fake news” I’ll be using in this and the next chapter.²⁸¹

Real News: The news, as defined above.

Fake News: An informational product meant to pass as *the* news, but which is not actually *the* news.

²⁸⁰Ladd (2012)

²⁸¹The term “fake news” is also used in a different sense, what I call the “rhetorical” sense, the discussion surrounding which has interesting implications for our beliefs about lay citizens and civic epistemology, but I will wait to address that issue until chapter five.

Testimony and Belief

All news, “real” or “fake,” is testimonial. When I use the term “testimony” I do not do so in the colloquial sense, as in “more or less official claims offered up by a witness to a police officer or a jury.” I mean it, rather, in the philosophical sense, as *any report or claim of fact made by an individual or institution to another or others*. According to this definition, testimony is a *genus* of which one *species* is news.

Testimony: Any report or claim of fact made by any agent or group of agents to any other agent or group of agents.

Whether we accept, reject or hold in abeyance a proposition brought to us through testimony is a question of *belief*. Colloquially, the word “belief” is used to describe a *feeling of conviction* with regard to a given proposition. In many cases of concern to epistemologists I believe that is a fine definition. However, this dissertation is about *politics*. Accordingly, I am only concerned with a phenomena insofar as they have *public implications*. I am not concerned with beliefs that merely reside in a person's mind yet do not affect their actions in any way. Beyond this, allowing the *mere* feeling of conviction in a proposition to define what counts as a “belief” begs irresolvable disagreements about whether or not a person who has manifested no external behavior in response to a proposition “really” believes that proposition. For these reasons, when I use the word “belief” in this chapter I will adopt the notion of belief championed by American pragmatists such as Charles Peirce and William James, which is *a proposition that serves as one's basis for action*.

Belief: A proposition that serves as one's basis for action.

We derive our beliefs about the world from a number of sources. The one I am particularly concerned with, both in this chapter and the dissertation as a whole, is testimony. Like other ways of gaining knowledge – introspection, perception, memory, and so on – testimony comes with its own set of advantages, disadvantages, inherent limits, pathologies, and modes of compensation. There is a branch of philosophy, the epistemology of testimony,

dedicated to exploring these issues, but it has not yet been applied to the way we as a public approach the news. To that task, and its implications for contemporary politics and liberal democratic theory, I turn in the rest of this and the next chapter.

Section Two: Doxastic Self-Rule as a Core Component of All Normative Epistemology

Civic Epistemology as an Underlying Concept in 21st Century Politics

A central normative concern for many epistemologists mirrors a central talking point in many discussions about the plight of 21st century democracy and the effect of emerging communications technologies on it. This is the question of *what we should believe*. From Plato's cave to Descartes' (supposed) stove to Wittgenstein's brash claim to have figured out everything and posthumous disavowal of the very possibility of humans ever doing so, epistemologists have been wringing their hands over the question of what (if anything) people ought to believe, why they ought to believe it, and how confidently they ought to hold whatever beliefs they form for millennia. The advent of the internet era, with its dramatic decrease in barriers to publication and increase in access to alternative sources of information, has prompted public-minded folk of all stripes to take up many of the same questions that have plagued epistemologists for all these years. Amid the deluge of information and onslaught of new and unfamiliar sources of it, the question of *what citizens ought to believe* is the focal topic of many political discussions, and serves as the implicit source of disagreement about many others, in the present day. I call the topic of what citizens ought to believe, and how they ought to go about adopting their beliefs, *civic epistemology*, and the various models of belief formation implicit (but seldom implicit) in commentators' praise and scorn for the epistemic behaviors of 21st century democratic citizens *models of civic epistemology*.

Civic Epistemology: The arena of discourse revolving around which beliefs democratic citizens ought to hold and/or how they should go about forming them.

Models of Civic Epistemology: Normative models of belief formation for democratic citizens.

I chose the word *civic* both because it implies a focus on the behaviors of *lay citizens* and their contribution to or detriment toward the public good, as well as because the term is historically associated with the concept of *virtue*. The discussions I am concerned with are discussions about the doxastic decisions of lay citizens and their contributions to or detriment toward the public good. Frequently, the subset of citizens whose doxastic behavior I aim to defend – the so-called “post-truthers” – are accused by Progressives of *epistemic civic vice*, that is, behaviors that are morally objectionable and/or detract from the public good. Conversely, the approach to belief formation Progressives endorse implies their views about which doxastic behaviors and habits ought to count as *epistemic civic virtues*.

While implicit in much contemporary discourse about democracy and 21st century communications technologies' impact on it, the term “civic epistemology” has as yet not been coined, nor has any other term yet been established capturing the notions of epistemic civic virtue and vice. Unfortunately, this is not just a terminological problem. It is instead an indicator of the glaring absence of rigorous examinations of the pressing and often extremely difficult questions begged by any normative model of belief formation – questions that have plagued epistemologists for thousands of years. In my view, the civic epistemology embraced by Progressives suffers from a number of problems that stem largely from their lack of a firm foundation in normative epistemology. Once we are clear about first principles, however, I believe it is quite clear that the model of routine deference to established epistemic authorities advocated by Progressives is incompatible with any compelling account of normative epistemology whatsoever, and especially any compelling model of civic epistemology applicable to liberal democracy.

Justified Belief as the Appropriate Standard for Normative Epistemology

Despite their many differences, nearly all normative epistemologists I am aware of agree on at least one thing. This is that we ought not to believe any old thing willy-nilly. Instead, we ought to exercise some degree of discretion when it comes to admitting propositions into our library of beliefs. Our beliefs, after all, serve as our basis of action and our actions affect both ourselves and others. Making sure our actions make our lives and the lives of others better, and avoiding the opposite, depend critically on our embracing sound beliefs and avoiding false ones.

It is depressingly easy to adopt misguided beliefs, beliefs which could wind up hurting ourselves and/or others, so we ought to make an effort to adopt accurate ones and avoid faulty ones. In other words, belief formation is a normative matter, and no convincing account of belief formation entails just haphazardly adopting any old belief.

But what *should* be the normative standard for belief formation? Ideally, we'd like people only to adopt those beliefs they know to be true. However, in human affairs there is the pesky problem of *chronic epistemic fallibility*. This is the simple and commonly acknowledged fact that we might always be wrong. No bullet-proof means has yet been discovered for guaranteeing that any given belief, or even any belief-formation process, will give us the truth. If anything, the opposite is the case: so embarrassingly pervasive is our tendency to err that the most reliable form of inquiry we have put together is to *assume* our best ideas are in error and to *incentivize* people to prove to everyone just how wrong those ideas are. This is the basic model around which our institutions of science, with its attempts to preserve free inquiry and bestowing of prestige and financial rewards to those who successfully overturn or improve upon precedent, is built. Chronic epistemic fallibility means that “truth,” insofar as it means “correspondence to reality,” is for all practical purposes irrelevant when it comes to determining whether one has adopted one's beliefs in an epistemically responsible manner. It is a desirable *aspirational ideal* to which the knowledge-seeker may aspire in order to help them do the best they can, but since they cannot really *know* whether they have attained that standard, it is not a useful way for us to assess whether they have acted responsibly.²⁸²

Instead of basing our assessments of people's epistemic performance on whether their resultant beliefs are *true*, it makes more sense to instead fall back on whether they are *justified*, given the limited resources available to mere mortals like us trying to navigate what William James famously called the “blooming, buzzing confusion” of reality.²⁸³ Instead of “Did we reach the truth?” then, the question we ought to ask in evaluating the epistemic performance of ourselves and others is, “Did we do the best we could?” That is, did we adopt our beliefs by using the best-available processes and the most reliable tools available to us in the situation? If we have, then our belief (or non-belief) is justified. If not, then not. In other words, because of

²⁸²This is not to say that we cannot *believe* the (real) truth. It's just to say that we cannot tell the difference between those things we *think* are true, but which are not, and those which we *think* are true, and really are.

²⁸³James (1981 [1890]), 462

chronic epistemic fallibilism, normative evaluations of belief formation ought to center around whether the process by which the belief came to be adopted is appropriate, and not whether the particular belief that is adopted in a given situation turns out, in the long run, to be true. This way of going about normative epistemology is called *process reliabilism* in the epistemology literature.²⁸⁴

Doxastic Self-Rule as an Essential Component of Normative Epistemology

Epistemologists disagree about precisely *which* behaviors we ought to employ in order to meet the demands of normative epistemology, but very many agree with what I have said so far about chronic epistemic fallibilism and its implications for making justified belief, rather than correspondence for reality, the proper normative standard for individual belief formation. I also believe they would agree on one more thing. This is that, whatever process a person must undergo in order to form justified beliefs, it oughtn't simply be an *arbitrary choice*. Nobody thinks writing down candidate beliefs on the various faces of a many-sided die, rolling it, and adopting the belief that appears on the face that winds up on top would be a praiseworthy process of belief formation. The problem with this situation isn't that we have it out for dice. Flipping a coin, drawing straws, casting lots, and blindfolding oneself and playing "pin the conviction on the proposition" would not count as epistemically responsible either. Why not condone belief-by-die-roll or belief-by-lot-casting? I believe the reason is (or ought to be) that epistemic responsibility requires the choice of whether or not to believe a proposition ought to involve *scrutiny in the light of the judgment of the individual knowledge-seeker*, or what I call *doxastic self-rule*. I believe our support for doxastic self-rule as an indispensable element of epistemic responsibility stems from two sets of reasons. First, doxastic self-rule is a necessary component of both deliberately improving our ideas over time and uncovering and resisting abuse by authorities. Second, doxastic self-rule embodies notions of self-rule that are a central element of most of our shared ideas about morality, particularly those of us who embrace liberalism and/or democracy. Let us examine each of these reasons in turn.

²⁸⁴For more on process reliabilism, complete with bibliography, see Goldman and Beddor (2016). For an influential contemporary account of process reliabilism, see Goldman (1986).

First, let's look at why I argue doxastic self-rule to be an indispensable component of improving our beliefs and uncovering and resisting abuse by authorities. Fallible though it may be (and it certainly is), I argue that such scrutiny in the light of one's own judgment is the *only* means of belief formation that is *not arbitrary* in the same sense as rolling a die or casting lots would be. That is, it is the only means we can employ to *deliberately seek to improve our beliefs and uncover and resist incompetence and/or abuse by authorities*, rather than leaving our epistemic and political fates to random chance. That non-arbitrary belief formation requires an act of judgment on the part of the individual, and that the regular use of such judgment is a more reliable means, over time, of helping us improve our beliefs can be seen by comparing it to some of the most common ways people go about forming their beliefs. Let us consider, for example, belief formation through *lived experience*. This serves as the basis for many of our beliefs, but only *some* of the beliefs we form through such experience have been scrutinized in the light of our own judgment. That distinction is normatively important because our brains are constantly in the habit of pattern-matching, but some of those patterns can be deceptive in ways that can be corrected by sound use of judgment. For instance, the everyday experience of many police officers in American cities exposes them to a disproportional amount of crime in Black neighborhoods. Left unchecked, the simple pattern-matching operation of the mind may, and I believe often does, lead many of these officers to associate Blackness with criminality. However, it is far more likely that those police officers' tendency to encounter criminality in Black neighborhoods stems from the endemic poverty, political and societal estrangement, and systemic discrimination the Black community has faced throughout American history than it does with the color of the people in those neighborhoods' skin. By scrutinizing the associations being automatically and constantly formed in their minds in the light of their judgment, police officers can resist the build-up of prejudice that might easily insinuate itself in the absence of such scrutiny, while to fail to subject those accumulating associations in the light of their judgment leaves them vulnerable to the build-up of (unjustified) prejudice. This sort of argument applies to many situations wherein the *active exercise of judgment can condition the beliefs formed through unconscious pattern-matching*.

Unconscious pattern-matching is not the only way we can form beliefs in a way that does not entail scrutiny in the lights of one's judgment. Another common way is by absorbing without

much question the ideas about *authority* handed down to us by the various communities to which we belong. This is routinely done in all areas of life. From politics to morality to religion to science, many people's beliefs about what is true stem from their ready acceptance of the claims given to them by those whose authority they have been taught by their community to accept, on the *mere* basis of having been taught to treat them as authorities. In many cases, those authorities deserve to be treated as such, but it is trivially easy to find examples, both past and present, where they do not. Sometimes this is because authorities make errors in good faith, as I assume to be the case for many of the members of the church who opposed the heliocentric model of the universe, burned countless books, and tortured and killed many perfectly decent people on grounds of heresy. Other times, the authorities embraced by a community – religious, political, or otherwise – do not deserve to be treated as such because they are morally corrupt and actively seek to manipulate that community in ways that serve the interest of the authorities at the community's direct expense, as I take to have been, and continue to be, so commonly the case in politics as to need no specific example. Either way, the habit of simply absorbing the prevailing views within the relevant community about which authorities deserve to be treated as such inhibits that community from discovering the authority's lack of legitimacy, no matter whether that lack of legitimacy stems from simple incompetence or from the authority's deliberate abuse of power. Resisting these outcomes requires the individuals within those communities to be willing to predicate their trust in those authorities on their *own judgment*, rather than simply simply absorbing the notions that prevail among their peers.

The examples of belief formation given above – belief formation through unconscious pattern-matching and through unquestioning acceptance of the beliefs handed down to one by one's community – are *arbitrary* from an epistemic point of view. That is, whether or not their use results in reliable belief formation is an *accident of chance*. I believe this is true of *all* beliefs formed on any basis other than the deliberate application of judgment by the belief-forming agent. *All* decisions and beliefs whose adoption is not conditioned on an active exercise of one's own best judgment are as arbitrary as if they rolled a die, cast lots, or flipped a coin. They are based on either the accident of one's idiosyncratic array of anecdotal experiences, or the the accident of being born in a given community, or something equally arbitrary in terms of the predictability with which it will result in sound beliefs. Such arbitrary processes *can*, and often

do, result in the adoption of sound beliefs. But no normative model of belief formation should endorse such arbitrary processes as the standard of belief formation to which we should aspire, because there is nothing normatively praiseworthy in *accidentally and incidentally* getting something right. We only approve of decision-making processes that are *arbitrary* from an epistemic perspective in situations where the results of those decisions are inconsequential – like outcomes in board games or decisions about which restaurant to go to – or for which our powers of judgment are assumed to be utterly useless. But in the former case it is the very *lack* of stakes that makes such an arbitrary process acceptable, and in the latter it is not that an arbitrary decision is *commendable* so much as it is that the direness of our circumstances makes even what would normally be a horrendous decision-making method as good as any other available to us. Even in the latter case, though, the decision to resort to a decision-making process whose outcome is arbitrary should *itself* be justified by the lights of one's own best judgment – that is, one's own best judgment *that* the stakes of the decision are too low to merit worrying about the outcome, or that our knowledge is so miniscule as to make our judgment on the matter at hand essentially worthless. Thus, for all decisions thought to have important real-world stakes, doxastic self-rule in the form of requiring one's beliefs and decisions to pass a test of scrutiny in the lights of one's own best judgment, is a fundamental component of normative epistemology.

The Principle of Doxastic Self-Rule

None of my endorsement of the consistent subjection of beliefs, and deference to authority, to scrutiny in the light of one's judgment implies that individuals who consult their judgment will *always* come to the right conclusions. On the contrary, I take it as inevitable that many people will make mistakes in judgment and that plenty of times these mistakes will be consequential. A police officer may reach racist views, people may embrace factually incorrect views of the world, and community members may wrongly decide that an untrustworthy group of authorities deserves their trust or that a trustworthy one does not, even through a sincere effort to base those conclusions on their own judgment. Our judgment is *chronically fallible*, and my argument does not imply anything to the contrary. It *does* imply, however, that over time and in the long run people's judgment is *good enough* that the regular employment of it will lead to the gradual improvement of individual thinkers' beliefs, and consequently the beliefs of any

community comprised of members regularly exercising doxastic self-rule, more reliably than if they adopted those beliefs arbitrarily – that is, on any basis *other* than their own judgment. Thus, my preference for doxastic self-rule implies a *modest optimism about human judgment*. I believe a similar modest optimism lies behind most of our general preference for people to adopt their beliefs, and choose which testimonies and testifiers to trust, on the basis of their own judgment rather than picking them at random or passively absorbing the accepted views of the communities to which they belong. Our judgment may be, and is, chronically fallible, but it is nevertheless our *best* means of choosing what to believe and whose testimony to trust. It is better for a police officer to deliberately subject their experiences to scrutiny in the light of their own judgment than to simply hope that the gathered associations that build up in their minds result in an accurate representation of society and the causes of criminality. It is better for people born into a given community to deliberately subject the behavior and claims of the authorities handed down to them by that community to such scrutiny than to simply hope that the community they happen to have been born into picked all of the “right” ones. In the long run and over time, the regular employment of judgment by individuals in these and other scenarios is a better guarantor of the steady improvement of our ideas and the discovery and overturning of illegitimate authority than the regular adoption of beliefs by such persons through any means that does not entail the active exercise of judgment on the part of individuals.

It is important to recognize that although I have framed doxastic self-rule as an *individual's* exercise of *their own* judgment, I have claimed that the *value* of doxastic self-rule manifests on *both* the level of the individual *and* of the collective. This applies both in terms of epistemology and in terms of power relations. Individually, the exercise of doxastic self-rule is lay citizens' only means of actually *evaluating* truth-claims. All other bases for the adoption of belief reduce to the *passive* or *accidental absorption* of prevailing notions, either within one's family or within some other community of influence, about what is true and/or which authorities' claims ought to be deferred to. Since every such community of influence is fallible and some, such as cults and abusive families, are actively corrupt, passive absorption of the notions that prevail within that community would doom each individual who practices it to the same flawed beliefs as those that characterize their community of influence and perpetual abuse by that community's authorities. As important as these individual-level effects are, however, the

collective effects of individuals' exercise of doxastic self-rule (or lack thereof) are even more profound. Again, this applies both epistemically and in terms of power relations. The improvement of our collective knowledge consists entirely in identifying and remedying errors that prevail in our current beliefs and in discovering new truths beyond those with which we are already acquainted. Both require an active exercise of judgment on the part of individuals and require those individuals to go beyond the ideas that prevail within their community of influence. Likewise, the collective goal of uncovering corruption and abuse by established authorities requires individual citizens to subject the actions and claims of those authorities to scrutiny in the light of their own judgment. Even though I have emphasized the need for *individuals* to exercise doxastic self-rule and scrutinize the testimony they encounter in the light of their own judgment, then, the *justification* of doxastic self-rule as a central component of civic epistemology stems from *both* its individual *and* its collective benefits, both in terms of epistemics and in terms of power relations.

I believe the case presented above is enough on its own to justify our embracing doxastic self-rule as a core component of epistemic responsibility. Yet, these do not exhaust our reasons for embracing that point of view. Further reasons can be found in the form of certain moral notions the vast majority of us share, which contribute to the appeal of democracy. The form of democracy most of us support – and certainly the form of it *I* support – encourages citizens to *think for themselves* come to their *own* conclusions about which individuals, issues, or parties to support, and put them into action in the form of a vote. In either case, whether support for individual candidates or policies or political parties, we see the proper basis of that decision as each individual citizen's *judgment*. The equal weight given to each vote would seem, too, to imply that we value each citizen's judgment equally. Thus, support for the widespread practice of doxastic self-rule is implicit in many of our views about how democracy ought to work. Our preference for doxastic self-rule can also be inferred from the shared features of political systems and power relations we detest, such as slavery, serfdom, and divine right, all of which appear to be objectionable at least partly because they violate the principle of doxastic self-rule. That is, by depriving certain individuals and groups any say in the laws and leadership by which they are ruled, they deny the relevance of those individuals and groups' judgment as a determinant of which laws to pass or which authorities deserve to be treated as legitimate and seek instead to

make them subservient to the will of others on an arbitrary basis, such as one's race, class, or bloodline.

In sum, I believe we have a multitude of reasons, both in terms of (a) their consequences for the improvement of beliefs and the discovery of illegitimate authority as well as (b) certain core normative ideals implicit in our embrace of democracy for supporting what I call the *principle of doxastic self-rule*, which is the notion that *people's decisions about what to believe ought to be conditional on their having scrutinized those beliefs in the light of their own best judgment*.

Principle of Doxastic Self-Rule: People's decisions about what to believe and whose testimony to trust ought to be predicated on their scrutiny of those beliefs and/or that testimony in the light of their own judgment.

Acknowledging Reasonable Limitations – Doxastic Self-Rule as an Aspirational Ideal

I believe the principle of doxastic self-rule to be the single most fundamental element of normative epistemology, both for individuals and for groups. For an individual to decline to subject their beliefs to scrutiny in the lights of their own judgment is to leave their beliefs and endorsements of authority to chance, thereby making themselves unduly susceptible to the adoption of misguided beliefs and support for illegitimate authority. Likewise, a society to be composed of individuals who routinely decline to subject their beliefs about the world in general and their notions of which authorities deserve their deference to scrutiny in the light of their own judgment will not easily improve the ideas that flourish within it and will very easily fall under the sway of incompetent and/or corrupt authorities. The reader should take note of the word *decline* in the prior two sentences. I chose that word deliberately. I take it as obvious that no individual can subject *all* their beliefs to scrutiny in the light of their own best judgment. We adopt far too many beliefs during our childhood, when we have neither the experience nor cognitive apparatus necessary to subject beliefs to much meaningful scrutiny, to go back through and examine every single one of them once we do. Nor can we, as adults, pay attention to every single one of the beliefs being passively formed in our minds through the constant process of pattern matching described in the example of the police officer above. Because of this, we cannot

hope to subject *all*, or perhaps even *most*, of our beliefs to scrutiny in the light of our own best judgment. If our normative standards for epistemic responsibility are to be reasonable (which I certainly believe they should), we cannot count as a failure anything less than the subjection of *every single* of their beliefs to the test of critical scrutiny in the lights of their judgment.

But this does not mean doxastic self-rule is *irrelevant* as a normative standard. It just means it is subject to the same reasonable limitations as any other normative standard of human behavior – that is, that we cannot reasonably demand perfection. Be that as it may, however, we *can* reasonably expect people to *make a general effort* to subject their beliefs to critical scrutiny in the light of their own judgment, and I believe any normative account of belief formation ought to encourage people to do so. Likewise, any persuasive normative account of belief formation should count it an epistemic failure for individuals to *deliberately decline* or *habitually neglect* to do so, and should oppose normative models of individual epistemology that embrace or encourage such habitual refusal or neglect.

Section Three: Trust, Testimony, and Epistemic Authority

The Role of Judgment in Justified Deference

It is worth pausing here to stress one thing this argument does *not* imply. My argument here does *not* mean it is always wrong to defer to the judgment of someone else. On the contrary, to refuse to *ever* defer to the judgment of someone else so would be catastrophic, crippling our ability to navigate this complicated world effectively. But my argument *does* mean that the decision of *whether* one should follow someone else's judgment rather than their own on a given matter *must itself be justified by an act of judgment*. Specifically, it should be based on a *comparative assessment* that seeks to evaluate whether a given testifier is *both* (a) trustworthy, in the sense that they are likely to be telling the truth, and (b) *more knowledgeable* on the matter at hand than oneself or the alternative testifiers available to oneself. To fail to predicate deference to the testimony of some authority on the exercise of this sort of judgment is to exercise *rote* or *automatic* epistemic deference, which, in turn, is to fail to exercise *doxastic self-rule* and to fail epistemically. This makes requiring trust in testimony to be justified through active scrutiny in

the light of one's own judgment a basic feature of normative epistemology, which brings us to the final principle I will articulate in this section. This is the *justified epistemic deference principle*.

Justified Epistemic Deference Principle: People ought only to defer to the claims of testifiers if they judge (a) their *own* knowledge to be inferior to that of the testifiers, (b) the testifiers to be *trustworthy*, and (c) the testifiers to be the *most credible* authorities available in the situation at hand.

The principle of doxastic self-rule gives priority to one's own best judgment, requiring one's beliefs – including beliefs adopted through trust in the testimonial claims of others – to be justified by its lights. Since one's own best judgment has priority, deference to others ought only be made in those cases where the lights of one's own best judgment lead one to believe that some other individual or entity is likely to both *know better* than oneself and *be a trustworthy authority* on which to base one's actions and/or beliefs. Simply judging them to know better is clearly not enough, for people who know better than the lay person may abuse deference to their authority as readily as anyone else – indeed, perhaps more easily.²⁸⁵ In order for that deference to be justified, then, a testifier must pass both an epistemic and a moral test in the eyes of the lay citizen in addition to the lay citizen judging their own knowledge on the matter at hand to be comparatively lacking.

Doxastic Self-Rule and the Justified Epistemic Deference Principle as Core Components of Liberal Democracy

I have so far defended my model of normative epistemology based on doxastic self-rule in general terms – that is, I think this viewpoint should be accepted simply on the basis of its

²⁸⁵Examples proliferate both in politics and in other realms. Political office holders are quite famous for using the asymmetry between their knowledge and that of the public for purposes of manipulation. In fact, that asymmetry is a plausible contributor in the U.S.'s decision to enter nearly every major armed conflict in the last century, from the sinking of the *Lusitania* to the Gulf of Tonkin incident to infants being “bayoneted” in Kuwait to weapons of mass destruction in Iraq. For just two examples outside the explicit realm of politics, psychologists in the 1950s and 1960s subjected patients who went to them for therapy to involuntary injections of LSD and extended periods in sensory deprivation chambers (Klein 2007), while recent spate of athletic trainers, prominently Larry Nassar, who in recent years have been found to have engaged in decades' worth of sexual abuse of patients who submitted to treatment on the grounds of those trainers' medical expertise.

being advantageous to the aims of improving beliefs and discovering and resisting illegitimate authority on both an individual and a collective level, as well as its being implied in the basic moral framework implicit in many of our embrace of democracy. The principle of doxastic self-rule leaves room for many decisions of trust that may be made in many epistemic domains by many different citizens. It leaves room for people who decide to place their trust in religious texts, news reports, scientific journals, government reports, the word of parents, and a panoply of other testimonial sources. In any of these domains, citizens' decisions to trust any of the sources of testimony mentioned *may* be justified or it *may not*. It is vital to emphasize that, consistent with the principle of doxastic self-rule, no outside observer can determine whether someone's decision to trust a given testifier is justified merely by pointing out the identity, social position, prestige, accolades, of or any other externally identifiable feature of the specific source of information providing the testimony. The reason is simple: doxastic self-rule requires individual lay citizens to subject those common signifiers of epistemic authority to scrutiny in the light of their own judgment, too. The strong preference for doxastic self-rule means that it is not the *mere* possession of such credentials or the *mere* holding of some particular social position that justifies deference to a given epistemic authority; instead, it is those authorities' having been *judged to be both credible and trustworthy in the eyes of the individual citizen*. If citizens have conditioned their deference to an epistemic authority on an active exercise of their own judgment, that deference is justified. If not, then not.

While I believe them to apply broadly, the strong preference for doxastic self-rule and the justified epistemic deference principle have a special place in liberalism, democracy, and liberal democracy. I have already made the argument connecting both doxastic self-rule and the demand for deference to authority to be justified to liberal democracy in chapter one, but the requirement that demands of *epistemic* deference be justified by the lights of citizens' own judgment is just as central to liberal democracy as these other two. One of the most important features of liberal democracy as I have defined it is its emphasis on *oversight*. The importance of oversight of authorities by the public is implicit in the core features of both democracy and liberalism and, *a fortiori*, their combination in liberal democracy. The defense I presented of democracy in chapter one implies that oversight is one of the most important reasons for supporting universal, equal suffrage among able-minded adults. Recall that I did not endorse the idea that the justification for

universal, equal suffrage is its ability to express the “public opinion” or the “general will.” The argument against thinking of the value of universal, equal suffrage as stemming from such grounds has, I think, been well prosecuted by any number of theorists.²⁸⁶ *While I think its ability to express “public opinion” or the “general will” is neither an accurate nor desirable basis for defending democracy, however, I do* think there is a significant benefit to requiring that political office holders gain and regularly retain the formal approval of the public in the form of a vote. One of the major values of universal, equal suffrage is that it exercises a restraining force on those office holders, what Richard Posner has called a “loose tether,”²⁸⁷ disincentivizing them from engaging in abuses of their power likely to draw too much of the public's ire. Meanwhile, liberalism's freedoms of speech and press serve as a check on authority by maintaining the conditions best suited to letting as wide a range of voices as possible publish and disseminate claims. As Mill argued long ago, guarantees of such freedoms are superfluous when it comes to opinions already approved of by the majority and by those in power, which presumably need no such protection.²⁸⁸ Instead, their main benefit is to facilitate the voices of *dissenting* groups, i.e. those disseminating opinions unlikely to be approved of the majority and/or those in power.²⁸⁹ Thus, (the best versions of) both liberalism and democracy emphasize the importance of maintaining public oversight of authority.

In both cases, it is critical to note that for this oversight to occur and to be effective requires more than just the proper constitutional setup. It also requires the adoption of the proper *disposition* toward authority – one of *vigilance* – on the part of the public. As Mill wrote in his *Considerations on Representative Government*, “if, from indolence, or carelessness, or cowardice, or want of public spirit, they are unequal to the exertions necessary for preserving it,” then that society lacks the necessary conditions for effective representative government, even if

286Lippmann's case against this sort of justification for democracy was made as well as any, first of all because the public does not usually know enough to choose policies well and second of all because simply marking a “cross on a piece of paper” is not an adequate means of capturing anyone's views about public policy (Lippmann 1925, chs. 4 & 5). Beyond this, social choice theorists have compellingly demonstrated that voting is a poor means of preference aggregation (e.g. Arrow 1950).

287Posner (2003)

288Mill (2007 [1859])

289This is the logic behind epistemologist David Coady's defense of rumors and conspiracy theories, terms whose usage he equates to “unofficial communications” (Coady 2011, 96) and “the respectable modern equivalent of ‘heresy’” (Ibid, 123), to which I am highly sympathetic.

the legal structures for such a government have been set in place.²⁹⁰ The constitutional enshrinement of freedoms of communication and universal, equal opportunity of suffrage among able-minded adults can attempt to foster the legal conditions that facilitate such vigilance, but the protection of those conditions would mean nothing were citizens to adopt a disposition of routine and uncritical deference to the very authorities those constitutional protections are meant to enable citizens to check. In order for those protections to mean anything in practice, citizens must *not* defer to the authorities presented to them as a matter of course, but instead must act on the basis of the justified deference principle, that is, insisting that any deference to would-be authorities be justified by the lights of their own judgment – specifically, their judgment that the authority in question is a *legitimate* authority, one whose declarations *ought* to be deferred to. Thus, while I believe the justified deference principle to be a normative behavioral standard that applies very broadly, I also believe it to be especially important to liberal democrats.

Trust in Testimony as the Basis for Most of Our Beliefs

I have now brought us from general considerations about normative epistemology to the specific topic of trust in testimony. Trust in testimony might seem like a niche topic, and indeed was treated as such by epistemologists for many years. However, in the past few decades probably due at least in part to the same societal developments that have caused so much worry about the epistemic well-being of democracy, this has begun to change.²⁹¹ Epistemologists have increasingly noticed that the use of the interrogative term *what* in the expression “What should we believe?” which serves as the foundational question for normative epistemology, is potentially deceiving. It directs our focus to the *propositional content* of the belief being formed rather than the *basis on which* that belief is formed. Asking the question “*What* should we believe” can easily lead us to suppose that the most important element of the belief-formation process, and the core normative component that drives the epistemic debates that rage in the internet era is people's disagreements about the *truth or falsehood of the specific propositions* different citizens adopt, or refuse to adopt, as beliefs. Yet, while the truth or falsehood of those

²⁹⁰Mill (1991 [1861]), 14-15

²⁹¹After being a central topic of concern for early modern thinkers like Locke, Reid and Hume, the topic of testimony was largely ignored by epistemologists until being resurrected by the debate between John Hardwig (1985, 1991) and Steve Fuller (1986). In the decades since, it bloomed into an entire, and increasingly populated, subfield, called “the epistemology of testimony.”

propositions, and the belief (or disbelief) of citizens in them, is surely important, in the vast majority of cases we do not decide what to believe primarily through an evaluation of the merits of that proposition *itself*. Nor should we. In a huge majority of cases, we do not have the requisite experience to judge the reliability of the proposition at hand on its own merits, nor the opportunity to acquire such experience to any appreciable degree. Instead, in the huge preponderance of cases – and certainly in the overwhelming majority of the cases that spawn heated discussions about the plight of democracy in the information age – the question of “what to believe” really reduces to the question of *whose claims to trust*.

The rest of this and the next chapter focus entirely on the epistemic questions surrounding *trust in testimony*. In them, I will argue against the model of civic epistemology implicitly embraced by Progressives and in favor of an alternative model of civic epistemology based on the widespread practice by citizens of doxastic self-rule and an embrace of the justified epistemic deference principle. I will base much of my argument on considerations about the epistemology of trust in testimony. Since it is not commonly realized just how many of our beliefs are based on decisions about whom to trust, the reader might at first think the epistemology of trust in testimony to be too much of a niche topic area to justify using it as the primary basis for an entire normative model of civic epistemology. However, I believe that its centrality to the overwhelming majority of our beliefs makes the epistemology of trust in testimony a topic with far more sweeping implications than it at first may seem. Though I frame my arguments in terms of *trust in testimony*, then, the reader should keep in mind that, since such a huge preponderance of our beliefs are formed on the basis of trust in testimony, the range of beliefs to which these arguments apply is quite massive.

Many contemporary political discussions revolve around the question of which claims and/or claimants lay citizens *ought* to trust in various situations. For example, the climate change debate centers around whether and how far citizens ought to trust the claims of the climate science community regarding the causes and likely effects of climate change. Similarly, much of the controversy surrounding the United Kingdom's decision to exit the European Union, or “Brexit,” consisted in those opposed to Brexit citing the arguments of economists and condemning those who voted in favor of epistemic malfeasance for failing to place sufficient trust in the claims of those economists. Likewise, citizens are divided in their opinions about the

news. Some view the legacy news media as the “fourth estate” and argue or imply that it ought to be largely trusted, while others view the legacy news media with deep suspicion and view the former group as excessively gullible. In all these examples, the central question is whether and how far citizens ought to place their trust in the claims of some *established epistemic authority* within a given domain – climate scientists within the domain of climate science, economists within the domain of economics, and legacy news journalists within the domain of the news. And, as we have already seen and will be demonstrated further, in all these examples, Progressives argue that citizens ought to trust the claims of the established epistemic authorities and accuse those who refuse to do so of epistemic malfeasance of a sort that threatens democracy itself.

For the rest of this chapter and the next, I will make the case against the civic epistemology embraced by Progressives and in favor of a civic epistemology based on doxastic self-rule on a series of considerations about the epistemology of trust in testimony. Progressives seem to have a sound grasp of the profound epistemic limitations faced by any individual lay citizens in the contemporary world, but their response to those limitations, which condemns doxastic self-rule and advocates routine deference to established epistemic authorities, is misguided and dangerous. By taking a closer look at the epistemic conditions that normally characterize lay citizens' decisions about whether and how far to trust would-be epistemic authorities, I will argue that the strong preference for doxastic self-rule and the justified deference principle are appropriate normative standards for civic epistemology even in the large, complex societies characteristic of 21st century democracies. Indeed, they are not only appropriate, but *necessary* in order to facilitate improvement of collective beliefs and oversight of authorities in the same ways described in the discussion of general normative epistemology above.

Reductionism and Non-Reductionism in the Epistemology of Testimony

The predominant question that has occupied scholars in the epistemology of testimony has been the debate between reductionists and non-reductionists.²⁹² The debate between

²⁹²For those interested in an introduction to this debate, see “Reductionism and the Distinctiveness of Testimonial Knowledge,” by Sanford C. Goldberg and “Testimony and Trustworthiness” by Keith Lehrer, as well as Jennifer Lackey's “It Takes Two to Tango: Beyond Reductionism and Non-Reductionism in the Epistemology of

reductionists and non-reductionists in the current literature revolves around the question of whether or not testimony ought to count as a “basic” source of beliefs alongside other “basic” sources widely acknowledged by epistemologists, such as perception, introspection and memory.²⁹³ According to reductionists, testimony should *not* count as a “basic” source of knowledge. That is, it should not be in the same category of other sorts of knowledge we acquire through *personal experience*, such as perception and memory. Reductionists tend to argue that beliefs adopted on the basis of testimony are only justified when they are based not *merely* on a given instance of testimony itself or even testimony that is *merely* corroborated only by other testimony, but instead can only be justified when that testimony is corroborated by “basic” knowledge of *some other type* – for example, by past personal experience with a testifier (which, in the epistemological literature, would reduce to a combination of perception and memory).

Non-reductionists counter that the position of reductionists profoundly underestimates our dependence on testimony to navigate day-to-day life. This objection was well articulated by 17th century Scottish philosopher Thomas Reid, who wrote in response to skeptics like Descartes and Locke that the “wise and beneficent Author of Nature...hath...implanted in our natures two principles that tally with each other.” These principles were, first, “a propensity to speak the truth, and to use the signs of language so as to convey our real sentiments” and, second, “a disposition to confide in the veracity of others.” Were we not so disposed, claimed Reid, “no proposition that is uttered in discourse would be believed, until it was examined and tried by reason; and most men would be unable to find reasons for believing the thousandth part of what is told them. Such distrust and incredulity would deprive us of the greatest benefits of society, and place us in a condition worse than savages.”²⁹⁴ This argument, cleansed of its now-unfashionable theistic and ethnocentric elements, remains the basic position of contemporary non-reductionists, who accordingly argue that it would be preposterous not to count the phenomenon to which we owe such a tremendous proportion of our beliefs as a “basic” source of knowledge.

Testimony,” which is something of an intervention intended to move beyond the debate. All are in the volume *The Epistemology of Testimony*, eds. Jennifer Lackey and Ernest Sosa (2006). For a moderate view of non-reductionism, see Lackey (2003). For a defense of strong reductionism, see Fricker (1995). For an attractive compromise hinging on the distinction between “justified” and “unjustified” knowledge, see Audi (1997)

293There is no consensus among epistemologists as to the exact specification of this list of “basic” sources, but these three are among the most widely acknowledged.

294Reid (1983), 93-95

Aspects of both these arguments are powerful. On the non-reductionist side, there is no denying that (a) we owe an unrepayable debt to uncorroborated testimony for giving us a tremendous bulk of our knowledge; (b) we could not by other epistemic resources acquire the thousandth part of that knowledge; (c) the great preponderance of human utterances are statements at least the speaker believes to be true; and (d) we would be hampered, indeed even crippled and fully incapable of navigating this world, were we from birth to withhold belief in testimony until it received support from some other type of knowledge. Yet, there is an appeal to the reductionist position as well. Be testimony a “basic” source of knowledge or not, it just seems *better*, if possible, to be able to corroborate testimony with our personal experience rather than to just rely on chains of testimony for justification all the way down. Are these positions really as incompatible as the “basic vs. non-basic” framing implies? I think not. If we frame the question not in terms of *whether* a belief ought to count as “basic” but instead of *how much confidence* we are justified in placing in the beliefs we adopt on the basis of testimony, I believe we can incorporate the insights of both reductionist and non-reductionists into a more sensible and practically useful framework for dealing with trust in testimony.

The strength of the non-reductionist position is its awareness of the vast swath of beliefs we hold, many of which appear to be extremely reliable, which we have adopted on the basis of testimony alone or at least on testimony that is corroborated only by other testimony. Meanwhile, the strength of the reductionist position is in insisting that it is *preferable* to get experiential corroboration for testimony before adopting it into one's beliefs, rather than relying only on testimony that is corroborated just by other testimony or, worse, testimony corroborated by nothing else at all. I believe there is an easy way to get the best of both worlds. This is by granting to non-reductionists that testimonial corroboration *can* be a strong enough warrant for the adoption of a belief to merit the label *justified* in the sense that it can *serve as a legitimate warrant for action*, while at the same time conceding to reductionists that we ought generally to *place greater confidence in* testimony that is supported by a non-testimonial warrant than in testimony supported by other testimony alone.

The Inscrutability of Testimony (by Itself) Principle

Not only is this a more pragmatically useful way of framing the reductionist vs. non-reductionist debate than the “basic vs. non-basic” debate that has tended to preoccupy epistemologists, but I believe it can also serve as the basis for a strong normative model for belief formation via testimony. First, though, I must answer what I anticipate to be a common objection. This is that I have over-emphasized *variety between* pieces of evidence in support of a belief and under-emphasized *quality of* that evidence. While I think considerations of the *quality of* evidence at our disposal to be an important consideration, I also believe that one of the peculiar pathologies of testimony is that it does not admit of the same ease of differentiation in evidentiary quality *by itself* as other knowledge types do. Take *reason*, for example. There is a *qualitative clarity* to the logic of the syllogism that $x + 0 - 0 = x$, such that even six and seven year-olds given a “conservation task” are able to discover its truth (in practice, that is, not algebraically) without ever being taught,²⁹⁵ whereas the reliability of the quadratic equation is likely to be more opaque to most comers. Or take another example, this time from *sensory perception*. There is a *qualitative clarity* to the pain of stubbing our toe on the leg of a coffee table that assures us of that coffee table's presence in a way that looking at the same coffee table through the window of a 747 at 15,000 feet cannot match. In the *clear* version of both these examples, we have every right, without bringing in any outside evidence whatsoever, to consider ourselves as having firmer grounds for confidence in their respective conclusions – that $x + 0 - 0 = x$ and/or that there is a coffee table with a leg that rests at foot level in our living room and will not budge easily upon being struck with a toe – than we do in the *ambiguous* versions. We can differentiate in some degree between *strong* and *weak qualitative grounds* for conclusions reached by syllogisms (reason) and encounters with coffee tables (perception) without ever having to resort to other types of knowledge at all.²⁹⁶

I believe we have nowhere near a similiar degree of intrinsic ability to reliably distinguish between ambiguous and unambiguous versions of *testimonial* evidence. Societies past and present have brimmed with liars, tricksters, scammers, propagandists and quacks. Many would

²⁹⁵This conservation task finding has been reproduced several times in the developmental psychology literature.

For a foundational study in the tradition, see Ames and Murray (1982).

²⁹⁶That said, we still can and often do err in both cases, so consilience still strengthens the justification of even our belief in foot-height coffee tables and six year-old syllogisms.

put politicians into that same group. Why have such peddlers of unreliable testimony proliferated throughout human history, while much less historical success has been had by anyone trying to claim that $2 + 2 = 0$ or that the everyday objects people stub their toes into are illusions? I believe one major reason is that, unlike our reasoning about basic arithmetic or our encounters with coffee tables, our ability to make distinctions between beliefs we ought to adopt on the basis of qualities of the testimony *itself* is severely limited. When liars, tricksters, scammers, propagandists and quacks (and politicians) make claims, we are all but helpless to tell by their testimony *alone* whether what they say is true or false. Instead, in order to evaluate their likelihood we must compare those claims against *something else* – the politician's track record, our estimation of the quack's credentials and the likely properties of snake oil, the payoff structure of the sales scheme we are invited to join, and so forth. The ability to tell, by external signs, when someone is lying is much touted in popular lore, but on closer inspection seems highly suspect. No particular behavior can be pointed out which correlates reliably with lying and nothing else. Even those behaviors often associated in our minds with lying, be it by politician or close acquaintance – failure to make eye contact, stuttering, fidgeting, or looking generally “squirrely” – are poor indicators of the quality of the testimony we encounter. Indeed, in cases these very phenomena may just as likely be indicators of *honesty* and *good character* as of scurrilousness and deceptiveness. For instance, they may indicate a testifier's being nervous *because* the testifier, being a person at once sensitive but possessing great character, takes himself to be telling a truth that will make present company uncomfortable but which needs to be said anyhow.

Such considerations make a case, based on folk psychology, that humans cannot dependably make distinctions between more and less reliable instances of testimony by simply analyzing the testimony *itself*. But we need not rely *solely* on folk psychological accounts of this sort to support the same conclusion. A meta-analysis on the scholarly literature on truth-discovery by Charles F. Bond, Jr. and Bella DePaulo showed that people are on average able to detect deception 54% of the time, a rate of success scarcely distinguishable from the insight that would have been granted by consulting the flip of a quarter.²⁹⁷ This ability does not seem to appreciably increase when we are dealing with those closest to us, as our ability to tell when

²⁹⁷Bond, Jr., and DePaulo (2006)

people we are familiar with are lying is little if any better than our ability to judge the veracity of the testimony of complete strangers.²⁹⁸ It may be true that despite these aggregate trends *some* individuals *can* tell with some reliability when a close acquaintance is fibbing. For example, they might know that their sister tends to look down when she is lying but not otherwise, or that their husband puts his hands in his pockets when he is attempting guile. But in those cases it is not the testimony of the individual making the claim *by itself* that justifies the conclusion they are lying but instead that instance of testimony assessed against the backdrop of listener's *antecedent knowledge* of the fibber that allows them to recognize the lie. Such “exceptions” merely prove the rule.

For all these reasons, based in both folk psychology that should make sense to all of us as well as social science research, I conclude that *the reliability of a given instance of testimony cannot be dependably assessed by evaluation of the testimony itself*. I call this the *inscrutability of testimony (by itself) principle*.

Inscrutability of Testimony (by Itself) Principle: The reliability of a given instance of testimony cannot be dependably assessed through a qualitative evaluation of the testimony itself.

From the *inscrutability of testimony principle* it follows that in order to gauge the worth of the testimony we've been given we must consult something else. It also strongly suggests that corroboration by *other* testimony that is not *itself* corroborated by non-testimonial means, is a *weak* warrant for confidence in a given testimony, for the absence of non-testimonial support leaves us as unequipped to gauge the reliability of the corroborating testimony as we are the testimony we are using it to corroborate. I will call testimony that cannot be corroborated by anything except other testimony “testimony *by itself*.”

The problems presented by the *inscrutability of testimony* will plague us right down the line, be the chain of testimony ever so long. Since we cannot dependably differentiate reliable versus unreliable instances of testimony through qualitative analysis of the testimony itself, it follows that testimony *by itself* is something of a precarious basis for belief in general, and

²⁹⁸Van Swol, Malhotra and Braun (2012)

therefore ought to be treated as a *limited* warrant of confidence. In fact, were what I have written so far all to be said on the matter, the conclusion would be even more damning. If there were no further considerations to take into account, it would seem that testimony *by itself* ought not to even serve as a *weak* warrant of confidence in belief, but instead give one *no* grounds for confidence in that belief beyond whatever grounds the knowledge-seeker might already have.

But as it turns out there *is* one further set of considerations to take into account – one that should, I believe, lead us to the conclusion that this attitude of deep skepticism toward beliefs adopted on the basis of testimony by itself would be incorrect. These considerations revolve around the difference between standards appropriate for belief formation in an *ideal* world and those appropriate for the world we *really* live in. As it turns out, in the *real* world, the widespread tendency for people to *tell the truth* as a matter of course and to *demonstrate sound judgment* in the adoption of most of their beliefs make testimony far more reliable than the account I have offered so far. Less hearteningly, that same real world also demands that we often act on the basis of beliefs of which we are far less than certain. I will argue that both considerations ought to make us more ready to act on the basis of testimony than we otherwise might.

Testimony and Decision-Making in the Real World

Given the argument I have made so far, some might be inclined to reach the reductionist-friendly conclusion that beliefs adopted on the basis of testimony *by itself* ought never merit enough of our confidence to truly count as being *justified*. I disagree with that conclusion. In a perfect world – one where no act need ever be taken or belief adopted without first meeting optimal epistemic standards – that might be right. But we, notoriously, do not live in a perfect world. Instead, we live in a world where many of the choices we make occur under *circumstances of tragedy*. What I mean by the term “circumstances of tragedy” here is that we are often forced to make decisions with real-world repercussions even though we know our grounds for making those decisions are less than perfect and that therefore we have no great assurance that our decision will result in good rather than harm. If we accept that our world is tragic in this way, we need to accept that people's beliefs can sometimes be *justified* – that is, used as a legitimate warrant for action – even if they do not meet *optimal* standards for belief formation. One reason I think beliefs adopted on the basis of testimony by itself *can* be justified,

then, is because in the real world, it is often better to act on the basis of imperfect information than to take no action at all.

Another aspect of the real world that ought to lead us to consider testimony *by itself* as meriting at least *some* confidence in a claim is that the overwhelming majority of testimony appears to be both obviously true and uttered in good faith by the testifier. Not only are we telling each other what we *believe* to be true the overwhelming majority of the time, but apparently we are also usually *correct in those beliefs*, at least insofar as “correct” means that the actions we take on the basis of those beliefs are effective in helping us pursue our desired ends. I believe these conclusions are demanded by any thoughtful reflection on societal affairs, political or otherwise. The world is far too enormous and complex for us to explore but the tiniest part of it by ourselves, yet we as individuals move and act within it with tremendous success. On many occasions, the projects we undergo – building skyscrapers and highways, managing hospitals and businesses, administering governmental programs – require not only an accurate view of the effects our actions are going to have on our part of the world but also an accurate understanding of the actions that are going to be taken by others. In a huge preponderance of cases, our beliefs in both respects are founded on nothing other than testimony for which we have no corroboration in personal experience – other people's claims about how the tool we're operating is likely to work, or where they'll be at 4:00, or which laws are in force that govern the program we're trying to help administer and what they mean for the paperwork we need to fill out. Given the *inscrutability of testimony* as defined above, it is impossible to imagine the resounding success with which humans generally coordinate their actions in such a world, oftentimes completing tasks requiring incredibly precise coordination, if speakers were not being *honest with* each other the vast majority of the time and were the claims they made not *largely correct*. In other words, human beings exhibit a general propensity to *adopt reliable beliefs* and to *tell (what they take to be) the truth*. Given the general human tendency to adopt reliable beliefs and to tell (what they take to be) the truth, it seems bizarre to insist that testimony by itself ought to be given *no* weight when deciding how much confidence to place in a given belief. On the contrary, though it might seem gullible at first, the position of non-reductionists who insist that our *a priori* disposition toward testimony ought be one of credulity is, I think, closer to the truth than the position of reductionists who argue in favor of an *a priori* disposition of skepticism.

Nevertheless, real-world considerations also necessitate that we resist making the adoption of such an *a priori* attitude of credulity, or what I prefer to call an *attitude of routine doxastic deference* to testimony the basis of any model of epistemic responsibility. For as much as people exhibit a *general* tendency to adopt reliable beliefs and tell (what they take to be) the truth, there are many exceptions. Even if *most* people's beliefs in *most* situations are quite reliable, our chronic fallibility means that, inevitably, *many* of our beliefs in *many* situations are not. Likewise, even if *most* people are honest *most* of the time, the advantages that are regularly available to individuals willing to lie for their own benefit and at the expense of others mean that *many* people are dishonest *much* of the time, and that nearly *everyone* is dishonest *sometimes*. Obviously, this is as true in politics as in any other arena. Any model of civic epistemology should try to mitigate against such dangers in a way that simply recommending for citizens to adopt an attitude of routine doxastic deference, either to testimony in general or that of any specific individual or group, clearly does not. This need is made all the more urgent once one considers that individuals who *are* willing to lie in order to pursue their own good at the expense of others are likely to do so all the more aggressively if they observe that their listeners routinely accept every claim they hear, and at least some individuals who would normally not be tempted to take advantage of the public in this way may be tempted to do so once they notice the ease with which they can get such a gullible audience to believe whatever they say. For these reasons, even if we acknowledge, as I believe we ought, that people are generally good at adopting reliable beliefs and usually tell the truth, we should still avoid models of civic epistemology that encourage the adoption by large segments of the public of an attitude of routine doxastic deference to testimony, either in general or by specific individuals or groups.

How *can* a model of civic epistemology guard against the dangers mentioned above? I argue the only available alternative to the sort of routine doxastic deference to individuals or institutions mentioned above, which I have argued to make citizens excessively vulnerable to deception and abuse, is encouraging them to exercise *doxastic self-rule*. That is, a model of civic epistemology ought to encourage citizens, to the greatest degree practical, to consistently evaluate *all* testimony they receive from *any* source *by the lights of their own judgment*, rather than adopting an attitude of routine doxastic deference either to testimony in general or to the testimony of specific individuals or groups. As argued above, the exercise of one's own judgment

is the *only* means by which an individual can seek to mitigate the threat of adopting mistaken beliefs and/or being deceived and manipulated by others, which are ever-present and apparently eternal dangers endemic to human affairs. The centrality of judgment to resisting false beliefs and discovering superior ones should be obvious. Were individuals not willing to resist testimonial claims of others on the basis of their own judgment, our knowledge would never improve. Instead, it would forever be plagued by the same mistakes and unable to grasp new truths beyond what is currently “known.” Likewise, were individuals not willing to resist the testimonial claims of others, humankind would be doomed to rest content with whatever current notions prevail about which individuals and institutions deserve to be treated as “authorities,” rather than uncovering the deficiencies of those who do not deserve that title and discovering the merit of those whose true desert of that title is as yet not widely known. Since those authorities would undoubtedly notice this routine acquiescence to their authority, this would inevitably result in steadily more abuse by those authorities of the trust the public places in them.

Section Three: Civic Epistemology – Epistemic Vigilance vs. Selective Skepticism

Chronic Fallibilism and Modest Optimism

It is vital to highlight the differences between the position I have articulated here and the sort of sweeping pessimism about both the general reliability of testimony and lay citizens' ability to evaluate its plausibility that has long reigned in intellectual circles as well as in our cultural mythology, and which currently infuses the ideas of Progressives. The approach I advocate toward testimony entails a certain *modest optimism* about lay citizens and their judgments about and use of testimony. The approach is *optimistic* because it sees most testifiers as making better-than-random judgments about the truth and honestly reporting their beliefs most of the time. At the same time, it is *modest* because, keeping in mind the principle of chronic fallibility, it recognizes that both the testimonies we encounter and our evaluations of them are bound to be flawed. It is also modest because it recommends that citizens exercise caution in the form of doxastic self-rule, seeking, as far as is practical, to condition their confidence in every testimonial claim they encounter on the basis of *their own judgment*, and in so doing make themselves less vulnerable to the easy adoption of their community's misconceptions and/or

abuse by authorities. Yet, even this doctrine of cautiousness toward the testifiers and testimony citizens encounter *itself* implies a modest optimism of another sort – a modest optimism about the judgment of lay citizens *in general*. For it implies citizens must be at least moderately good – that is, better than random – at determining which testifiers and testimonies are reliable and which are not, as well as which would-be authorities are credible and/or trustworthy.

This characterization of my view as optimistic might seem at odds with my earlier and ongoing emphasis about the importance of keeping in mind our *chronic fallibility*, the fact that any single one of our ideas could be fundamentally misguided, and many of them could probably be improved. But I believe that, to the contrary, my emphasis of the importance of chronic fallibility is fully compatible with the optimism I have expressed about human judgment. There is no contradiction between claiming that *in general* humans are pretty good at adopting reliable beliefs, testimonial and otherwise, and acknowledging at the same time that *any particular one* of our beliefs may be fundamentally misguided and that nearly all of them could probably be improved. Likewise, there is no contradiction between claiming that *in general* people's testimony is both factually reliable and honest, yet that we still ought to subject *each* of their claims, so far as we can practically manage, to scrutiny in the light of our own judgment. In both cases, conditions of chronic fallibility, both epistemic and moral, demand that we limit our confidence in, and maintain an open mind to the possibility of corrections of, or improvements upon, any particular proposition. Yet, crucially, both cases still manage to maintain a *fundamental optimism about human judgment* because the regular and vigilant exercise of human judgment is the very *means by which* our beliefs are seen to improve. Our chronic fallibility means we might be wrong about any particular belief, but the general reliability of our judgment means that by regularly scrutinizing those beliefs in the light of our own judgment we can slowly and steadily improve them over time. Likewise, the general reliability of our judgment means that we can fend off at least some of the dangers of being deceived by erroneous or deliberately misleading testimony. In other words, in both cases *the improvement of our beliefs depends critically upon the exercise of human judgment, and therefore implies a certain modest optimism about it*. The argument presented here is, I believe, the same combination of an appreciation for the scope of chronic fallibility and modest optimism about human judgment found in the second chapter of John Stuart Mill's *On Liberty*. That mixture of positions is well captured in the passage

where Mill explains that proper explanation of the “preponderance among mankind of rational opinions” owes *not* to any “inherent force of the human understanding,” but rather to the fact that our “errors are corrigible” and capable of being rectified through “discussion and experience.”²⁹⁹

(Selective) Pessimism Skepticism in the Progressive (and Conventional) View

Some of the views I articulated in the previous sub-section are shared, at least circumstantially, by Progressives. For example, Progressives seem to endorse an attitude of epistemic vigilance toward testimony, as evidenced in the deep suspicion they advocate toward claims forwarded on social media³⁰⁰ and the “blogosphere,”³⁰¹ as well as forms of testimony such as rumor³⁰² and conspiracy theory.³⁰³ The frequency with which the latter two terms are used among the public as shorthand dismissals of the claims described by them illustrates that Progressives' suspicion toward rumors and conspiracy theories is shared by many contemporary citizens, and survey results routinely demonstrate that their negative opinions about the reliability of information found on social media are likewise largely shared.³⁰⁴ These pervasive views reflect certain prevailing notions about testimony that might seem in some ways to recommend a vigilant approach to belief formation. For instance, the familiar “telephone game” is taught to many children as an object lesson whose apparent intent is to illustrate just how undependable testimony is as a vehicle for information. Insofar as these aspects of the Progressive view, and these prevailing notions, seem to recommend epistemic vigilance, they may seem to mirror my own views about civic epistemology. However, I believe those views to be so far removed from my own in important respects as to be antithetical to them, for the skepticism they endorse implies not a *modest optimism* about lay citizens' judgment but in fact in most cases a *deep pessimism* about the judgment and morals of lay citizens themselves. Moreover, the sort of epistemic vigilance the views of Progressives and our collective cultural

²⁹⁹Mill (2007 [1859]), 80

³⁰⁰E.g. Vosoughi et al (2017), Lewandowsky et al (2017), House of Commons Report (2018)

³⁰¹Goldman (2010)

³⁰²Sunstein (2014)

³⁰³Sunstein and Vermeule (2009), Sunstein (2016)

³⁰⁴Indeed, PEW researchers Mark Jurkowitz and Amy Mitchell call Republicans' and Democrats' shared distrust of social media sites as a rare “oasis of bipartisanship,” with fewer than 20% of respondents indicating that they trusted any one of the six websites mentioned as a reliable source of information about political and election news (Jurkowitz and Mitchell 2020).

inheritance seem to recommend is not *regular* and *sweeping* vigilance toward *all* testimonial claims but instead *situational* and *highly selective* form of vigilance toward *certain* testimonial claims – if, indeed, such views can be said to advocate any sort of epistemic vigilance on the part of lay citizens at all.

It may not seem immediately evident that the views attributed above to Progressives and our cultural inheritance express a deep pessimism about lay citizens' judgment, but I believe a close inspection of both those views themselves and others held by these groups should make this quite clear. Who is it, after all, that is responsible for the majority of the content on social media and the “blogosphere,” and what is it that differentiates social media and “blogosphere,” which Progressives distrust, from the legacy news media, which Progressives view as largely trustworthy and indeed indispensable to democracy? The major differentiating factor between these forms of media is that the information posted on social media is largely *uncurated by authorities* and shared, by and large, by *lay citizens* while that published in the legacy news media is *actively curated by authorities* and is written and edited by *professional journalists*. Progressives who see social media as a threat to democracy often cite its lower barriers to publication vis-a-vis that of the legacy news media as the primary *reason* we should suppose its information to be less trustworthy.³⁰⁵ In other words, it is precisely *the fact that* it is so easy for any member of the lay public to access and share information on social media that makes it such an unreliable source of information, and vice-versa for the legacy news. The implication would seem necessarily to be that either the honesty, or the judgment, or both, of the lay public is/are deeply and regularly unreliable, otherwise it is difficult to explain why we should suppose social media to be a democracy-threatening reservoir of misinformation and fake news.

A similar degree of pessimism toward the public (and friendliness toward established authorities) is implicit in Progressives' embrace of our widely shared cultural suspicion of rumors, which are viewed with such skepticism that the word “rumor” itself is treated as a synonym for an unreliable claim. That this skepticism implies a deep pessimism about the public and its willingness and ability to reliably make judgments about and pass on testimonial claims can be inferred by looking at which testimonial claims are *not* rumors. We do not call legacy news stories, press releases by Apple and other companies, or announcements by members of the

305E.g. Deb, Donahue, and Glaisyer (2017)

defense department about an ongoing foreign conflict “rumors.” What do all these non-rumor forms of testimony have in common? They are *official reports* by some *institution* widely viewed to be an *authority* on the matter at hand. By contrast, a “rumor” is an *unofficial claim* passed on through word-of-mouth or other informal means, such as social media, which is *not confirmed by any institution* that is widely viewed as an authority on the matter at hand. In other words, what *makes* a rumor a “rumor” is precisely that it is the *lay public* that is primarily responsible for passing it on and that it is not officially confirmed by some established, authorizing institution. The implications of Progressives' (and conventional) widely shared distrust of “rumors” thus mirror those of Progressives' (and conventional) widely shared distrust of social media. In both cases, that distrust betrays a deep implicit pessimism about the judgment and/or honesty of the lay public, a lay public that is viewed largely as either incapable of or unwilling to identify and pass on accurate information.

Riding alongside this Progressive (and conventional) *mistrust* of social media and rumors is an implicit (sometimes explicit) *veneration* of and *credulity toward* established epistemic authorities. Contrast Progressives' scorn for social media with their views of the legacy news media. The former they depict as a cesspool of misinformation and fake news that is a threat to democracy, while the latter they regularly call by the honorary title “the fourth estate” and characterize as a democratically indispensable institution whose employees work “in the service of truth.”³⁰⁶ Similar implications derive straightforwardly from a consideration of Progressives' (and conventional) treatment of rumors. In this case, readiness to believe rumors – i.e. informally transmitted claims passed on by lay persons but not confirmed by an established institutional authority – is treated as a potential threat to democracy. In both examples, it is *readiness* to trust the information passed on by fellow lay citizens on social media that is routinely treated by Progressives as a potential threat to democracy, but when it comes to the legacy news media the exact opposite is the case: now, it is a *refusal* to trust that is the threat to democracy. The clear implication would seem to be that lay citizens' judgment is deeply suspect, to the point where regular dependence on it threatens democracy. Meanwhile, the legacy news media is an indispensable safeguard for democracy, to the point where the refusal to trust it is a threat to democracy. This further substantiates the case that Progressives hold a deeply pessimistic view

306UNESCO (2018), 9

of lay citizens' judgment since a large and long-increasing proportion of citizens reports that it does *not* place much trust the legacy news media, which Progressives frequently interpret as a democracy-endangering catastrophe of judgment on the part of the public.³⁰⁷

So we can see that the sort of “vigilance” endorsed by Progressives is quite different from the one I have endorsed. The position I endorse is founded on a *modest optimism* about lay citizens in that it views citizens as both being largely honest and possessing generally reliable judgment. As such, it advocates for citizens to try, as far as practical, to exercise doxastic self-rule over *all* their beliefs, and so recommends that citizens subject *all* claims to the *same* doxastic standard, which is the light of their own judgment. It is *that* act of judgment, and not their having been “confirmed” by this or that institutional authority, that determines whether their belief has been responsibly formed, that is, formed in the way that best contributes, in the long run, to the growth of their own and collective knowledge while holding out the best prospects for uncovering and resisting illegitimate authority. In contrast, the Progressive (and conventional) view implies a *deep pessimism* about lay citizens' honesty and/or judgment. The Progressive (and conventional) view encourages citizens to be *selective* in their scrutiny of testimonial claims. It explicitly prescribes for citizens to view rumors they encounter and the claims they find on social media with a default attitude of suspicion and skepticism while not only encouraging them to adopt an attitude of credulity toward established institutional authorities such as those in the legacy news media, certain government bodies, and other knowledge-disseminating entities that enjoy a current reputation for being epistemic authorities, but routinely disparaging and condemning as threats to democracy those who do not. Whereas the civic epistemology I advocate is one of *consistent epistemic vigilance, including* vigilance toward the claims of established institutional authorities, the civic epistemology embraced by Progressives is one of *selective skepticism*, encouraging citizens to adopt a default attitude of skepticism toward forms of communication driven by lay citizens and a default attitude of credulity toward the claims of Progressives' favored authorities.

307Brenan (2020)

Selective Skepticism as a Threat to Democracy – Fake News as a Case Study

I believe the attitude of *selective skepticism* advocated by Progressives to be misguided for a number of reasons. By defining which authorities ought to be trusted, and advocating for a default disposition of trust in those authorities prior to and outside of any act of judgment on the part of individual citizens, selective skepticism undermines our best prospects for the growth of knowledge and the discovery of and resistance to illegitimate authority. Indeed, it cultivates a disposition of credulity toward established authority, suspicion of the claims of fellow lay citizens, and doxastic timidity that *actively encourages* abuse by authorities and *actively discourages* citizens from using the very faculty that is most likely to help them navigate the biggest challenges to the epistemic well-being of democracy in the 21st century – the faculty of their own judgment. That Progressives are excessively pessimistic about the judgment of lay citizens, and that the attitude of credulity they advocate toward their favored authorities is likely to cause democratic problems, can be demonstrated using one of the most-favored examples of what is wrong with the post-truth movement among Progressives, the example of fake news.

The term “fake news” emerged sensationally in the collective consciousness in the 2016 election campaign between Donald Trump and Hillary Clinton and remains for many the primary example of the post-truth malaise. The belief that we have a dire fake news problem and that it is indicative of the depths to which the epistemic well-being of democracy has fallen is widespread, and those who share this view often associate fake news with social media³⁰⁸ and attribute its (supposed) influence to high levels of epistemic vice among the public³⁰⁹ and low levels of trust in the legacy news media.³¹⁰ Because of the public's supposed ineptness in dealing with fake news and the putative danger it poses to democracy, Progressives have stepped up demands for regulation of the informational environment by authorities, be they “fact-checkers” employed by social media companies themselves or governmental authorities. And, indeed, numerous countries, including France and Germany, have passed legislation aimed at combating this supposed threat to democracy.

Yet, neither the *existence* of a threat to democracy by fake news nor the public's inability or unwillingness to combat fake news appears to be justified by the empirical literature. Both in

308DeB, Donahue, and Glaisyer (2017), Schiffrin (2017), Waszak et al (2018)

309Resnick (2017), Persily (2017), Vosoughi et al (2017)

310Schiffrin (2017), Persily (2017)

the United States and in Europe, empirical investigations into the scope of fake news have consistently found it to constitute an extremely small proportion of the contents shared via those media.³¹¹ Such studies routinely find fake news to be dwarfed by content deemed to be more reliable by the researchers conducting the studies, whose perspective on which sources ought to count as reliable closely mirrors that of Progressives. Democratic citizens' apparently tendency to easily resist fake news suggests Progressives' characterizations of them as riddled with epistemic vice that so inhibits them from rendering sound judgments as to make fake news a threat to democracy globally³¹² are excessively pessimistic, to say the least. Even in the informational environment in which, as Progressives are wont to remind us, fake news can travel farther and faster than ever before, lay citizens' judgment appears to be largely up to the task of fending off whatever threat fake news might pose to democracy.

Equally suggestive is the other most consistent finding in the literature, mentioned briefly in Chapter Two, which is that the single greatest predictor of susceptibility to fake news is age, with the oldest citizens many times more likely than the youngest to share fake news. While even the oldest groups still do not appear to share much fake news, those who grew up with the internet as an everyday part of life in their formative years demonstrate consistently better habits of link-sharing than their older counterparts. I believe this too suggests support for my arguments about relying on lay citizens' judgment as the basis for civic epistemology, rather than cultivating in them a disposition of deference to the “right” authorities, for two reasons. First, younger citizens' superiority in navigating fake news suggests that the ability to differentiate between real and fake news online can be *learned* through the sorts of ordinary experiences citizens are likely to have by virtue of simply growing up in the 21st century's online-dominated informational environment. Second, older citizens' comparatively higher susceptibility to fake news may plausibly stem from their having grown up during an era when citizens were routinely taught to look for “reliable sources” by looking for institutional markers that marked sources as members of the legacy news media. It is precisely those sorts of markers the most popular fake news sites traded on in order to gain the credulity of their viewers. For example, the most-shared fake news article in the 2016 election campaign was originally published by a (now-defunct) hoax news site

311Alcott and Gentzkow (2016), Fletcher et al (2017), Guess et al (2017), Nyhan et al (2018)

312Deb, Donahue, and Glaisyer (2017), Persily (2017), Schiffrin (2017)

called *WTOE 5 News*, a title that clearly trades on the old convention of television and radio stations adopting four-letter call signs starting in K or W. Many of the other sites that published comparatively popular fake news stories in 2016 used similar tactics.³¹³ *It is plausible that older generations' comparatively high susceptibility to fake news* was amplified by their tendency to trust any claim made by an entity bearing such indicators of belonging to the legacy news media. Insofar as that is the case, the default disposition of credulity toward the legacy news media advocated by Progressives may easily have been a *contributing cause of* citizens' susceptibility to fake news. Not only does the empirical literature on fake news seem to suggest the depth of Progressives' pessimism about lay citizens' judgment when it comes to telling reliable from unreliable news online is excessive, then, but the attitude of routine deference to the claims of the legacy news media may *itself* contribute to the very sorts of epistemic problems Progressives fear from fake news.

But the problems with Progressivism's deep pessimism about lay citizens' judgment and the default attitude of credulity they advocate toward authorities go far beyond their contribution to the apparently quite minor influence of fake news on the public. This disposition is now actively contributing to the growing tendency for major social media platforms and some governments to exercise censorship over the information available to citizens online. For instance, political communications scholar Kathleen Hall Jamieson recently praised French President Emmanuel Macron for suppressing certain information, whose characterization by the French government as “fake news” Jamieson accepts without elaboration, about Macron on the eve of the 2017 French Presidential election. The ability to engage in such suppression, argues Jamieson, makes France less “vulnerable” to the threat of fake news online than the United States, which is hamstrung by its strong protections of free speech embodied in the First Amendment.³¹⁴ *While not all Progressives so openly condone* governmental regulation of information online, the same combination of deep pessimism about the lay public and credulity toward their favored authorities naturally leads them to condone the same general approach of appealing to *some* set of authorities to fix the supposedly catastrophic problems in our

³¹³*The Denver Guardian, World News Daily Report, The Burrard Street Journal, and abcnews.com.co* all appear on Craig Silverman's list of the publishers of the 20 most popular fake news stories from August to November 2016 (Silverman 2016).

³¹⁴Jamieson (2018), 11-12

informational environment. Oftentimes this is done with a certain amount of hand-wringing, as Progressives sense the tension between their views and the ideals of liberal democracy. Aviv Ovadya, for instance, writes that we are “careening toward an infopocalypse – a catastrophic failure of the marketplace of ideas,” and although he describes himself as recommending “investment in human judgment,” in reality all his solutions invoke an appeal to some set of authorities, such as fact-checkers who will establish “certain baseline truths” for the public and “authenticity stamps” by browsers and platforms designed to “affect the psychology of belief formation, as merely labeling something as forged may not be sufficient.”³¹⁵ Likewise, in an article whose title poses the query whether democracy can “survive the internet,” Nathaniel Persily follows up his apparently strong liberal-democratic claim that democracy depends on “both the ability and the will of the voters to base their political judgments on facts” with the quick addendum “or at least on strong intermediary institutions that can channel decision making within a broad range of democratic alternatives.”³¹⁶ Given the strongly negative depiction of the public and its judgment implicit in Persily's account of the present informational environment, it is not hard to guess which two of these sources he sees as the more reliable guarantor of democracy's well-being. As a result of the increasing sway of this pessimistic view of the judgment of the lay public and the health of our informational environment, social media companies and governments around the world have stepped up efforts at censorship, framing it as a matter of defending democracy and preserving public health in the face of what the public health community has dubbed the “infodemic.”³¹⁷

Such policies are not exclusively the impetus of power-grabbing politicians and corporate tycoons running roughshod over a resistant people. On the contrary, in the case of large platforms like Facebook and Google, they were only adopted after resistant executives of these companies, citing ideals of free speech, were bludgeoned into submission by outraged politicians and many members of the public who, touting Progressive notions about the supposed “threat” posed to democracy by fake news and implying the same deep pessimism about the judgment of lay citizens, looked to those giant corporations to save the public, apparently from itself. That approach, I argue, is largely consistent with the deep pessimism of the lay public's judgment and

315Ovadya (2018), 43-45

316Persily (2017), 72

317WHO (2021)

credulity toward favored established authorities implicit in the Progressive position. Given the profundity of the doubts about lay citizens' judgment inherent in the view of them as entirely incapable of dealing with fake news, it seems natural that people convinced fake news is a threat to democracy would seek solutions to that threat that do not rely on lay citizens' judgment, and given their veneration of certain favored established authorities, it seems natural that it is to those authorities people would turn in the event of an “infodemic.” The biggest problem with the Progressive position is not necessarily that it is internally *inconsistent* (although it *is* inconsistent with liberal democratic ideals). It is just that it is *misguided*, both in terms of its characterization of the public and its public policy recommendations. There is no evidence that suggests fake news is even much of a problem, let alone that it “threatens democracy” or constitutes an “infopocalypse.” Thus, the case of fake news offers no support for Progressives' deeply pessimistic conclusions about the (un)reliability of the lay public's judgment. Worse, it encourages citizens to adopt an attitude of distrust toward their fellows that *itself* undermines democracy in a number of ways. It undermines the very *idea* of democracy being a good political system, because such deep doubts about the public naturally engender deep doubts about democracy itself. It may also undermine democracy because engendering a disposition of mutual distrust in citizens, especially one infused with the idea that one's fellow citizens are endangering the democracy one cherishes, plausibly undermines the basic “ties that bind” together any community engaged in a collective enterprise.³¹⁸ Finally, it undermines democracy because it discourages citizens from remaining open to the “unofficial” channels of communication most likely to discover and muster resistance to illegitimate authority³¹⁹ and contribute to the growth of knowledge.³²⁰

Just as Progressives' disposition of *deep pessimism* about the judgment of the lay public is likely to undermine liberal democracy, the disposition of *credulity* toward certain favored authorities embraced by Progressives is likely to do the same. First, it is not obvious why the public ought to suppose the collection of epistemic authorities favored by Progressives deserves to be treated as such. Again, this can be illustrated in the case of fake news, whose implicit

318The importance of such ties has historically been emphasized by democratic theorists like de Tocqueville and Mill (Miztal 2001), and has been emphasized in recent decades by scholars studying “social capital,” e.g. Putnam (2000) and Newton (2001).

319Coady (2011), chs. 4-6

320Mill (2007 [1859])

opposite is “real news,” or the legacy news media, which Progressives largely view as a bastion of truth, trust in which is vital to the well-being of democracy. It is not at all clear, however, why the public should share that laudatory view of the legacy news media rather than that expressed by their fellow citizens, two-thirds of whom agree that the reports of legacy news media are typically infused with a political, commercial, or other sort of bias and many others of whom believe the legacy news media to frequently focus on “exaggerated and sensationalized stories” and to be guilty of “inaccuracy and low standards,” and accordingly report low levels of trust in those reports.³²¹ Scholarship in communications, too, gives us reasons to doubt whether ready attributions of trust in legacy news by the lay public would be beneficial to democracy. Numerous scholars have presented empirically backed arguments that legacy news is excessively negative³²² and focused on sensational and non-representative events, or “outliers.”³²³ The legacy news' fixation on such negative and sensational events may explain why those who watch the news most regularly exhibit the greatest tendency to over-estimate the frequency of mass shootings, terrorism, and national crime rates.³²⁴ Beyond this, critics have long argued that the legacy news' historical reliance on a selection of disproportionately white and affluent journalists who undergo uniform training and rely on a homogenous and predictable staple of information sources and information-gathering routines is likely to result in a skewed picture of the world and make them easily manipulable by politicians and PR firms.³²⁵ Likewise, media critics in the heyday of professional journalism saw their close proximity to, and dependence on, those in power as producing a corrupting force on those journalists, making them too uncritical of the individuals they developed relationships with.³²⁶ I believe all these critiques have at least some merit, and while I do not think they ought to lead us to conclude that the legacy news media has no beneficial role to play in democracy, I *do* believe it suggests citizens ought to subject the news

321Newman and Fletcher (2017). Another study, conducted by Gallup in tandem with the Knight Foundation, found that, when asked for their reasons for mistrusting certain news organizations, 42% of Americans gave reasons falling under the category “Biased/Slanted/Unfair reporting,” while an additional 23% gave reasons falling into the category “One-sided/One point of view/Incomplete/Unbalanced/Not whole story” (“Indicators of News Media Trust” 2018).

322For a run-down of the extensive communications literature on negativity bias in the legacy media, see Soroka (2012).

323For a review of this also-extensive literature, see Kleemans et al (2009).

324E.g. ANES (2016)

325See, e.g., Boorstin (1987 [1961]), Sparrow (1999), Cook (2005), Bennett (2016)

326Crouse (1973)

to scrutiny in the same way Progressives want them to scrutinize the information they encounter on social media. That critical disposition would give citizens the best chance to fend off whatever reports of the legacy news media are actively misleading and draw the appropriate conclusions from those that are not. Meanwhile, the attitude of decreased vigilance suggested by Progressives subjects citizens to whatever flaws may be present in the reports of the legacy news media, deterring democracy from reaching its full potential for both knowledge acquisition and discovery of and resistance to illegitimate authority.

But the need for epistemic vigilance toward authorities goes beyond simply making the most of the legacy news as it currently stands. It is also a standing necessity for deterring the growth of corruption in *any* set of authorities, epistemic or otherwise. Even if the legacy news media truly *is at present* the bastion of democracy-preserving truth Progressives imply it to be, citizens *still* need to maintain vigilance toward it, otherwise some set of authorities in charge of the legacy news media will eventually notice the easy credulity with which their claims are accepted and seek to take advantage of it. Epistemologist David Coady makes this argument well in defending the role of rumor, or “unofficial communication,” in democratic societies. It is worth quoting the relevant passage in full:

“It may be that in an ideal society official information would carry an epistemic authority such that it would almost always be rational to believe it. But that is not our society, nor, I suspect, is it any society that has ever been or ever will be. What is more, if such a society were to come into existence, it seems likely that it would be unstable, since it would likely lead to complacency about officialdom that would be exploitable by officials hoping to manipulate public opinion to advance their interests. To the extent that the view that we should place our trust in official information rather than rumor gains widespread acceptance, official information will be less subject to scrutiny and, as a result, less likely to be true.”³²⁷

³²⁷Coady (2011), 99

Evidence of Complacency and/or Corruption in the Real News

Coady's argument rings true with regard to *all* would-be epistemic authorities in a democratic society, including those favored by Progressives. Even if a given authority is not corrupt at present, cultivating a default disposition of credulity toward, and deference to, that authority is likely to result in its becoming complacent and/or corrupt over time. And indeed this appears to be the trajectory taken by the legacy news media in the eyes of many Americans, who viewed it with high esteem in the early days of journalistic professionalism but have increasingly come to view it with suspicion and distrust. It is not at all obvious that this view is fundamentally misguided. The recent history of the legacy news is fraught with examples of what may plausibly be interpreted as evidence of complacency and/or corruption. This evidence includes the legacy news media's failure to unearth – or apparently even *investigate* – the truth about whether or not the Saddam Hussein regime was manufacturing weapons of mass destruction in the lead-up to and early stages of the Iraq war and legacy media icon Dan Rather's use of forged documents to impugn the military record of then-presidential candidate George W. Bush less than two months before the 2004 election date. It includes *CNN* contributor Donna Brazile having leaked questions to be used in the upcoming presidential debates to the Hillary Clinton campaign in the 2016 election cycle. Progressives frequently characterize Americans who view the legacy news media as left-biased and largely representing the views of those in power as driven by partisanship,³²⁸ lump them into the same broad category as loonies and conspiracy theorists,³²⁹ and insist that their suspicions toward the legacy news make them more vulnerable to fake news and misinformation³³⁰ and, of course, threatens democracy.³³¹ Yet, all of the examples given above can reasonably be interpreted as evidence that Americans have legitimate reasons to doubt whether the legacy news media is routinely giving them a fair, complete, and accurate view of the world.

There is another thing all the examples listed above have in common. In each case, where the legacy news media failed to uncover the truth, the unprofessionalized back channels of the internet came to the rescue. While the legacy news media was uniformly regurgitating the Bush

328E.g. PEW (2019), Owen (2020)

329Bauer and Nadler (2018)

330Satariano (2020)

331Hetherington and Ladd (2020)

administration's story about weapons of mass destruction, the blogosphere kept alive criticism and questioning.³³² It was on the blogosphere, too, that lay citizens first aired the case that the documents used by Rather had been forged.³³³ Wikileaks alerted the public to the fact that a member of the legacy news channel in charge of the upcoming town hall meeting had given one side advance notice of the contents to be used in that debate.³³⁴ The internet has empowered citizens to spot the potential flaws in legacy media coverage in other respects, too. Independent “citizen journalists” such as Tim Pool utilized emergent 21st century communications technologies to stream their own coverage of the Occupy Wall Street movement online, giving interested co-citizens more extensive coverage and a different point of view compared with that of the legacy news.³³⁵ Twenty-first century communications technologies are almost certainly responsible for the onset of the Black Lives Matter movement, as mobile technologies for the first time ever gave citizens in Black communities a means of capturing the unprovoked violence of police officers toward members of their community and disseminating direct evidence of that violence to fellow citizens who do not routinely suffer such treatment by police. Likewise, when the legacy news media focused on the sensational and violent elements of the protests that emerged after the videos capturing police brutality were released, depicting those movements as riots, citizens could turn for a different perspective to live-streamers who covered different aspects of the protests than those focused on by the legacy news media. I believe all these to be not just cases where the legacy news media failed democracy, then, but also cases where the resources made available by 21st century technologies aided it. To the extent that citizens in these and other similar examples were to adopt the attitude of *selective skepticism* advocated by Progressives, readily accepting the images and stories passed on to them by the legacy news media and viewing the information provided them through social media and the blogosphere

332See Loewenstein (2008) for numerous examples.

333The investigation was spearheaded by the writers at powerlineblog.com, a point that is little-emphasized in legacy news media accounts (Hinderaker 2004).

334Jamieson (2016)

335Stelter (2011). In fairness, the Occupy movement *did* eventually get mainstream media coverage, but a great deal of that coverage was a meta-analysis *by* the media *of* the media's coverage of the Occupy movement. Perhaps unsurprisingly, the verdict was self-exculpatory, with Eric Randall of *The Atlantic* concluding an acerbic critique of what he clearly viewed as the whiny Occupy protestors by noting that, since the NYPD police chief showed no signs of attempting to curb police brutality, the movement may soon have a “new angle” for basking in a “spotlight in which to complain that the spotlight isn't bright enough” (Randall 2011).

with the eye of skepticism and suspicion encouraged by Progressives, the civic epistemology embraced by Progressives harms democracy.

The Epistemic Vigilance Approach to Fake News

The reader should *not* interpret my argument here as implying that citizens ought to embrace an attitude of selective skepticism in the *opposite* direction from that endorsed by Progressives, viewing the reports of the legacy news media with skepticism and suspicion while lending easy credence to the reports they find on the blogosphere and social media. The reports of bloggers and “citizen journalists” may as easily be biased and flawed as those of the legacy news media, so it is just as important for citizens to subject their claims to scrutiny in the light of their own judgment as those of the legacy news media. My problem with the Progressive view is *not* that it wrongly advocates epistemic caution when it comes to the claims found on blogs and social media. Instead, it is that it *selectively* singles out those sources of information as unreliable sources prior to any act of judgment and simultaneously pre-defines the legacy news media as a reliable source in which citizens need to trust in order for democracy to survive. In so doing, Progressives seem to encourage citizens into a disposition of *excessive credulity* toward the legacy news media and one of *excessive skepticism* toward the alternative sources of information made available by 21st century communications technologies. This makes citizens excessively vulnerable to the flaws and biases of the legacy news media while simultaneously deterring them from making best use of the primary source of information they might use to help them discover those flaws, recognize those biases, and temper the beliefs they adopt partly on the basis of legacy news reports in the light of alternative perspectives. Instead of encouraging citizens to adopt a *selective* attitude of skepticism toward sources *predefined* as reliable/unreliable, we ought to be encouraging citizen to practice *epistemic vigilance* toward *all* sources of information. That is, we ought to encourage them to, to the greatest extent practical, actively scrutinize *all* testimony they encounter, from *any* source, in the light of their own judgment, and condition their trust in that testimony on that judgment.

Of course, in order to advocate this approach, we have to have at least a certain degree of optimism in citizens' judgment. Is such optimism justified? I believe so. Let's return to the topic of fake news. I have already argued that the apparent unwillingness of the vast majority of

citizens to share fake news justifies a *prima facie* suspicion that citizens' judgment is already, without any changes whatsoever, reliable enough to relieve most of our worries about fake news. This impressionistic argument is supported by the early research done by those seeking to find out what makes citizens susceptible to fake news. Whereas the Progressive view sees susceptibility to fake news as stemming from extreme partisanship, cognitive biases, and other epistemic vices that cast doubt on the very ability of citizens to make sound judgments in the online environment, the research of cognitive psychologists tells a different story. Such research is painting an increasingly clear picture that, in the words of Gordon Pennycook and David Rand, those citizens who *do* fall for fake news do so “because they *fail* to think; not because they think in a motivated or identity-protective way.”³³⁶ Pennycook and Rand found that inducing participants to engage in cognitive reflection – i.e. to *actively* scrutinize the claims they encountered in the light of their judgment – improved participants' truth-discernment *both* for committed partisans and for centrists, regardless of whether the claim they read aligned with their ideology or not.³³⁷ Elsewhere, Bago and colleagues report that activating heightened levels of attention on the part of participants decreased subjects' belief in false headlines but did not decrease belief in true ones.³³⁸ Both these articles also indicate that participants in general, whether prompted to undergo cognitive reflection or not, are significantly more likely to believe true stories than false ones. Likewise, after comparing the habits of those adept at differentiating reliable from unreliable claims online, Wineburg and colleagues concluded that instead of the “checklists” that proliferate in online guides for navigating fake news, which “focus on the most easily manipulated surface features of websites,” such as a title similar to those used by the legacy news media or the use of a URL ending in “.org,” education seeking to get citizens to avoid fake news and embrace sound information should focus on teaching them to (1) “read laterally,” consulting a variety of alternative points-of-view both with regard to the story at hand with regard to the source being consulted, (2) exercise “click restraint” rather than “mindlessly clicking on the first or second result,” and (3) teaching students to “use *Wikipedia* wisely.”³³⁹ The first two studies both vindicate citizens' powers of judgment in the sense that they show that

336Pennycook and Rand (2019), 48

337Ibid.

338Bago et al (2020)

339McGrew et al (2017)

citizens are *already* pretty good at resisting fake news *and* that their powers of judgment, *if activated*, increase that ability. The third study shows that people who are adept at identifying phony information online are good at it because they *use* their judgment *in tandem with the resources made available by 21st century communications technologies* to exercise discernment. These findings give support, in the specific domain of fake news, to the argument of Kevin Arceneaux and Ryan J. Vander Wielen that rather than viewing citizens as hopelessly ruled by their emotions and cognitive biases in the way Progressives – following an unfortunate trend in social science – tend to do, we ought to see citizens as capable of improving their cognitive performance by actively engaging in reflection and “taming their intuitions,” and encourage them to do so.³⁴⁰

In other words, research in the cognitive psychology of fake news suggests support for *modest optimism* about lay citizens' capacities for exercising sound judgment in the 21st century informational environment, as well as for the prospective benefits of encouraging citizens to practice *epistemic vigilance* by exercising *doxastic self-rule* while navigating that informational environment. Meanwhile, the empirical research on the actual spread of fake news appears strongly at odds with Progressives' bleak depiction of the online informational environment as an epistemic quagmire and, implicitly, of lay citizens as incapable of or unwilling to identify and resist fake news. Moreover, Pennycook and Rand's finding that the chief cause of fake news is citizens *failing* to think, rather than their thinking being fundamentally distorted, ought to arouse even more concern about the attitude of selective skepticism embraced by Progressives, for that attitude encourages citizens to only engage in active scrutiny of testimonial claims in *certain* situations, and implicitly not to be as active in their scrutiny of others. If the problem with fake news (to whatever extent there *is* a problem with fake news) is primarily one of citizens *failing* to think, we should resist any model of civic epistemology that implies that they only need to actively scrutinize the testimony they encounter in certain well-defined situations, as I believe the Progressive model does. Such a disposition does not properly prepare citizens to epistemically engage with the dynamic and quickly-evolving informational landscape of the 21st century.

340Arceneaux and Vander Wielen (2017)

Conclusion

In Chapter One, I argued that a certain modest optimism about lay citizens' judgment and a strong preference for those citizens to exercise doxastic self-rule is an integral part of any liberal democracy worth defending. In Chapters Two and Three, I argued that the regular exercise of doxastic self-rule by citizens is likely to aid the epistemic well-being of democracy, and that a broad consideration of historical trends, both recent and distant, gives us every reason to be modestly optimistic about the fitness of lay citizens' judgment for contributing to the growth of collective knowledge and facilitating good political outcomes over time, even if sometimes citizens exercising doxastic self-rule do so in a way that seems misguided from the standpoint of individual knowledge-seeking. In this chapter, I defended the same combination of modest optimism about lay citizens' judgment and the encouragement of the assertive exercise of doxastic self-rule from another angle, this time starting from the standpoint of the individual knowledge-seeker/citizen and arguing that the strong preference for doxastic self-rule is a necessary component of *any* convincing account of civic epistemology. I have argued that doxastic self-rule is the best means, both for individuals and for communities, to *improve their beliefs* and to *discover and resist illegitimate authority*. In this chapter I have given special emphasis to the topic of *trust in testimony*, which underlies most contemporary conversations about the epistemic well-being of democracy. I have argued that Progressives are wrong, not necessarily in each of their specific choices about *which* epistemic authorities to trust, but in their propensity to accuse anyone who reaches any other conclusions of epistemic civic vice. Implicit in the Progressive position is a view of civic epistemology that implies an unjustifiably deep pessimism about lay citizens' judgment, an excessive veneration of established epistemic authorities, and the encouragement of an attitude of *selective skepticism* likely to lead democratic citizens to be credulous toward established authorities and suspicious of their own fellow citizens in a way that is harmful to democracy. I used the the much talked-about issue of fake news, and its implied opposite the “real news,” to illustrate the ways Progressives' views of our current informational environment, and the citizens that dwell within it, are excessively pessimistic, and the approach to civic epistemology they implicitly advocate on the basis of that pessimism is likely to cause more problems than it solves. Instead of the deep pessimism and selective skepticism advocated by Progressives, I argued, we ought to view the judgment of citizens with a

certain modest optimism and encourage them to practice a civic epistemology of *epistemic vigilance*, regularly and actively scrutinizing *all* testimonial claims in the light of their own judgment and basing their attributions of trust not on a source's professionalism, officialdom, or other external signifiers, but rather on their *own* assessment, conditioned on the lights of their own judgment, about whether that testifier is likely to be both a *credible* and *trustworthy* source in the situation at hand.

In the next chapter, I will continue this discussion based in the normative epistemology of trust in testimony, taking a closer look at the tools available for would-be knowers trying to decide whether and how far to believe the testimony of sources whose reliability they cannot corroborate through any appreciable degree of personal experience, again using the topic of the news as a working example. I will argue that Progressives' treatment of the legacy news media and “fact-checkers,” as well as the censorial policies they often advocate, illustrate their over-readiness to attribute to authorities the status of “expertise” and to demand that their fellows ought to defer to them as a matter of civic duty. I will argue that, while *everybody* in most cases has non-negligible epistemic resources for determining whether and how far to trust a given testifier in a given situation, *nobody* ought to be counted as an “expert” in the skill of deciding which testifiers ought to be trusted. This, I will argue, gives lay citizens even more reason to rely on their own judgment in deciding whether and how far to trust the testifiers they encounter.

Chapter Five: Consilience, Confidence, and Justified Deference

In the last chapter, I argued against the *selectively skeptical* model of civic epistemology implicitly embraced by Progressives because it undermines citizens' best means for improving their knowledge and discovering deception and abuse by authorities, both in individual and collective terms. Against this Progressive view, I defended a model of civic epistemology that encourages citizens to exercise *epistemic vigilance* by scrutinizing the claims of testifiers to active scrutiny in the light of their own judgment. In making that case, I connected considerations about the epistemology of testimony with the concept of *news*, broadly conceived. In this chapter, I will continue this vein of thought. This time, I will object to the notion of “expertise” implicit in Progressives' treatment of news and the attitude of routine deference to the claims of legacy journalists they advocate on its basis. I will argue that Progressives' conviction that they themselves or any group of “fact-checkers” they can identify is much better placed, epistemically, to tell which news is reliable and which is not is highly suspect, as it implies a sort of expertise, *expertise about news*, which we have strong fundamental epistemic reasons to believe that *no one* can possess, and even if someone did, that expertise could not be demonstrated to the public in a way that would justify the expectation of routine deference to that individual or group's judgments. Since the expectation of deference cannot be *justified* to the public, it is incompatible with the ideals of liberal democracy, which, as argued in Chapter One, demand that all deference to authority be justifiable to those citizens whose deference is demanded. These epistemic concerns give lay citizens even more reason, above and beyond the already-compelling *practical* reasons I introduced in chapter one, to oppose both model of civic epistemology supported by Progressives and the censorial policies toward the management of misinformation and fake news of which many of them approve.

The argument will proceed as follows. In Section One, I will articulate a normative model of *justified doxastic confidence* based on the idea of *consilience*, or convergence upon a single hypothesis from a diverse array of evidence. I will argue that the framework of consilience

highlights both the potential and limitations of citizens when it comes to deciding what sources of testimony to believe, using the news as a working illustration. The next three sections will use this same framework and same working illustration to take an in-depth look at one of what I believe to be three of the best epistemic resources available to *anyone* trying to determine whether and how far to believe the news, lay citizen or (so-called) “expert.” These three resources are, in order of appearance, (1) consultation of a testimonial source's *reputation*, (2) consultation of one's *worldview*, and (3) the seeking out of *perspectival diversity* between testimonial sources. In Section Five I will lay out what I take to be the key implications of this discussion. The most important of these conclusions is that, while lay citizens need not feel *powerless* to reach justified beliefs on the basis of these tools, their *confidence* in the conclusions reached by *anyone* working under such conditions – themselves *and* any would-be “expert” or “fact-checker” – ought to be *limited*. This, I will argue, undermines the notion of “expertise” implicit in the notion that “fact-checkers” ought to be deferred to as epistemic authorities about the news, and supports the notion that citizens ought to decide the matter for themselves.

Section One: Consilience and Confidence

Consilience as a Warrant of Doxastic Confidence

In the last chapter I introduced the debate between reductionists and non-reductionists in the epistemology of testimony. The key difference between these two camps hinges on their views about *corroboration through personal experience*. Reductionists see the ability to corroborate testimonial claims through one's personal experience as playing the key role in determining whether believing a given testimony or testifier is justified, whereas non-reductionists do not. While agreeing with non-reductionists' claims that (a) *most* of our beliefs are adopted on the basis of testimony we cannot corroborate through personal experience to any appreciable degree and (b) many of those beliefs seem to be perfectly reliable and therefore ought to count as *justified*, I also agreed with reductionists' insistence that it is *better* to be able to corroborate the contents of a testimony or the credibility and/or trustworthiness of a testifier than to rely on testimony simpliciter. In that chapter, I did not pause to explain *why* personal experience has such epistemic value, instead relying mainly on what I anticipate to be the widely

shared intuition that it does. But I believe my argument against the Progressive pessimistic and deferential model of civic epistemology and in favor of one that is more optimistic about lay citizens' judgment and encourages them to wield it more assertively can be augmented by returning to that question.

Why *is* it better to be able to corroborate the contents of a testimony or reliability of a testifier is a better warrant of confidence in a given testimony through personal experience than not to be able to do so? I believe there are several reasons. One is that we have a kind of ability to be aware of our own mental states – our alertness, our goodness of intentions, the messages sent through our senses by our environment, etc. – that we don't have when it comes to the mental states of others. While these resources are not infallible, they *do* give us the capacity to assess the likely reliability of our own conclusions about the events we have experienced in ways that we cannot assess the quality of the conclusions other people reach about experiences they have had but we have not. I think this is what lies behind most of our inherent tendency to have more confidence adopting a belief adopted in response to have something *demonstrated before our eyes* than in response to just being *told that something is the case*. If we ourselves witnessed some phenomenon, we have access to the memory, not just of our perception of that phenomenon, but of the mental state we were in when we had that experience. We have access, for example, to how we *felt* when we had that experience – how sleepy we were, how excited an emotional state we were in, etc. – and how *vividly* or *clearly* the phenomenon we [think we] witnessed impressed itself upon us, in a way that we simply do not have access to the mental states of those who relay such phenomena to us through testimony. As with all of our epistemic tools, of course, such memories of our own mental states are imperfect, our proper use of them depends on our own self-honesty and sound judgment, and some are bound to use them better than others. Still, having access to this sort of potentially powerful set of epistemic resources or not is a significant differentiating factor, and part of the reason we are justified, *ceteris paribus*, in attaching more confidence to beliefs adopted on the basis of testimony for which we have some sort of corroboration through personal experience than on testimony for which we have no such corroboration.

However, there is a second, less intuitive reason I believe we are justified in having more confidence in beliefs in testimony that is corroborated by personal experience than testimony that

is not. This second reason has to do with the concept of *consilience*. Consilience is the convergence of multiple, systematically unconnected (at least, as best we can tell) pieces of evidence upon a single hypothesis, and is widely recognized to be a valuable indicator of a hypothesis' strength.³⁴¹ The intuition that drives this valuation is the notion of *resilience through diversity* familiar to ecologists, economists, nutritionists, and any number of other disciplines. The principle of consilience I want to support here is what I call the *general principle of consilience*.

General Principle of Consilience: The greater the diversity of independent (from our perspective) pieces of evidence that all suggest the same conclusion, the stronger, *ceteris paribus*, our confidence in that conclusion ought to be.

My embrace of this principle is based on a principle I have already introduced and invoked a number of times. This is the *assumption of chronic fallibility*, the idea, popularized by John Stuart Mill and commonly acknowledged in the abstract but, as Mill himself noted, just as commonly overlooked in practice,³⁴² that *all* human knowledge is susceptible to error. Mill emphasized our chronic fallibility in terms of knowledge's *acquisition*, but it is equally important to acknowledge its fallibility in *transmission*. The assumption of chronic fallibility means that we can never be certain either that we have *discovered* the truth or that we have *reported or received a report containing* the truth. Another way of expressing this thought is to note that we cannot tell for certain which things we currently believe to be true are *actually* true and which ones we *merely think* to be true.

It is commonly accepted that the assumption of chronic fallibility means we cannot tell whether any *individual* piece of evidence *actually proves* that a given proposition is true. But its implications are far more profound than that. For the assumption of chronic fallibility applies not only to our *beliefs themselves*, but to each and every one of our *grounds for belief*, as well. I mean the term “grounds for belief” very broadly. Such grounds might be particular piece of

³⁴¹For a defense of consilience as a sign of epistemic reliability, see McGrew (2003).

³⁴²“[W]hile everyone well knows himself to be fallible, few think it necessary to take any precautions against their own fallibility, or admit the supposition that any opinion of which they feel very certain may be one of the examples of the error to which they acknowledge themselves to be liable” (Mill 2007 [1859], 78).

evidence, but they might also be an entire research tradition or methodology. Because any of our grounds for belief might be flawed, we ought necessarily to consider conclusions supported by only one or very few grounds to be *precarious* compared to those beliefs that gather support from a diversity of grounds, for the simple reason that in the former case any flaw endemic to our single grounds undermines our *entire basis* for belief in *all* conclusions derived therefrom. Conversely, we ought to place more confidence in a conclusion to the degree that they are derived from a *diversity of grounds* for belief, where, again, the term “diversity” applies very broadly. A diversity of grounds for belief can mean, for example, a diversity of testifiers coming from different backgrounds and perspectives yet reporting a very similar thing; a diversity of evidence from different research traditions pointing toward a single conclusion; and/or a diversity of considerations taken from different “basic” sources of knowledge, such as reason, introspection, testimony, and so on, all seeming to point to a given belief. These considerations suggest that, all else equal, *the more diverse the array of grounds that subsidize a given belief, the more confidence we are justified in attaching to that belief*. This is because grounds of belief that are very different in kind are less likely to share a common set of flaws compared to those that are highly similar, and therefore even if one or a few of those sources is subject to error we may well still have strong grounds for adopting that belief.

As we will see in a moment, the ease with which a claim from a single or narrow range of sources can spread through testimony makes it very common for a preponderance of support to arise on the basis of a single grounds, and therefore consilience is an especially important desideratum to seek when it comes to deciding how much confidence to place in testimonial claims. But while it has *special* value when it comes to testimony, I believe the argument that, due to chronic fallibilism, doxastic confidence ought to be heavily determined by the degree of consilience enjoyed by a proposition, applies very broadly. It applies to research done by particle physicists just as much as it does to a New Year's resolutioner sifting through Google results in an effort to decide which diet is most likely to lose him those 20 pounds. Most importantly for the topic at hand, it applies to democratic citizens trying to figure out how much confidence they ought to attach to the beliefs they form on the basis of the testimonies to which they have access in the Information Age. In making that decision, I believe the rule that ought to guide their belief-formation process is the *consilience as a warrant of confidence principle*:

Consilience as a Warrant of Confidence Principle: The greater our diversity of grounds for adopting a given belief, the more confidence, *ceteris paribus*, we are justified in attaching to that belief.

Strong Versus Weak Consilience for Testimony

Though I believe consilience to be a prime desideratum for *all* our beliefs, here I am most concerned with testimony in particular. When it comes to testimony, I argue that consilience serves as a metric of justified confidence in two senses, which I will call a *strong* and *weak* sense. To the degree a proposition we encounter through testimony is corroborated by *personal experience*, we have *strong* consilience for that proposition. Conversely, to the degree a proposition we encounter through testimony is corroborated only by *chains of other testimonies*, we have only *weak* consilience for that proposition. To see why I call the former *strong* consilience and the latter *weak* consilience, let's look at an example. I am an avid gardener. I especially love growing things from seed. Now, seeds vary both in type (the species and variety of plant they will become) and in quality (germination rate and health of plant that will grow from the seed). It is important to gardeners who wish to grow things from seed to be confident that the seeds we purchase are *both* the right type and of high quality. In the era of the internet there are many sources of seeds available to gardeners like me, from enormous corporations to hobbyists discoverable on sites like Ebay and Etsy. What epistemic tools are available to internet-era gardeners to help them make sure they secure seed that is both of the right type and of high quality? One resource is *personal experience* with specific seed suppliers. Another is *recommendations by personal acquaintances*, such as friends whose gardening ability we admire. Still another is *online reviews and ratings* of seed suppliers and even specific varieties they supply. On the framework I have offered, the first two are examples of *strong* consilience and the latter an example of *weak* consilience.

My argument is that a gardener who has personal experience with the seeds supplied by a given company, or who is close to someone with such experience whose judgment and honesty they trust, has better resources for evaluating the reliability of that company's claims about the seeds it supplies than someone who must solely rely on the ratings and reviews of unknown others found online. The reason is that in the former two cases, though to varying degrees, the

gardener can put their personal experience to good effect in attempting to corroborate the *credibility* and *trustworthiness* of the reviewer. The greater the degree to which the gardener can use their experience to evaluate the performance of the seed supplier – the more extensive their experience with that company, the more gardening acquaintances they know who have used that company, and so forth – the *stronger* their consilience for the claims made by that reviewer. However, in many cases gardeners do not have access to this kind of *strong* consilience and must instead rely on ratings and reviews written by people they do not know. In those cases, gardeners can still get consilience for the various claims we encounter about our laptop purchase, but that consilience itself comes only in the form of *other testimonies* by people whose judgment and trustworthiness they have little ability to evaluate. Such corroboration by unknown others is better than nothing, but it is still *weak* compared to corroboration by personal experience because the gardener has little ability to evaluate the judgment and trustworthiness of those on whose testimony she is forced to rely. The more heavily gardeners must rely on *only* such testimony, and the less they can bring their personal experience to bear in evaluating the reliability of those testifiers, the *weaker* the consilience for their claims.

Now, I take it as obvious that, though I have called it “weak,” we can still put chains of testimony of the sort I have described to very good use in making decisions like buying seeds for our gardens. Many of those of us who live in the 21st century do so on a regular basis, relying on reviews and ratings posted by entirely unknown others who purchased (or claim to have purchased) a given product, tried out (or claim to have tried out) a recipe, been treated by (or claim to have been treated by) a doctor, etc., to help us make good decisions in an enormous and bewildering world. I call this kind of consilience “weak” not because it is *useless* or *suspect* but simply because, all else equal, it is a weaker basis for confidence *compared to corroboration through personal experience*. The intuition behind this argument can be isolated by considering cases in which where the testimonies we encounter about a given product or a given service provider online conflict with our personal experience. Many of us would still decline to source seeds from a company they personally had poor experience with in the past, even if Google Ratings gave them 4.7 stars based on input from 3,872 voters. Many of us, too, would prefer to go to a doctor with whom we, or at least a close friend whose judgment and standards we trust, have had extensive positive experience than to rely solely reviews from unknown persons posted

online. In these and similar situations, testimonial corroboration in the form of reviews and ratings by unknown others *can* be and often *is* a good means of making decisions; however, our preference for relying on personal experience even in such cases demonstrates that, at least when it comes to important decisions, many of us prefer to act on the basis of propositions for which we can find some sort of corroboration through personal experience rather than to simply rely on chains of testimony by largely unknown others.

Undoubtedly, some of this preference stems from overconfidence in ourselves or in those toward which we feel great affection. However, I believe our preference for placing our trust in claims that can be corroborated by personal experience also has strong epistemic grounds. One of these is the ability to qualitatively assess our own mental states, as argued above. Another, however, is that *personal experience is a fundamentally different form of knowledge³⁴³ than testimony*. In terms commonly used by epistemologists, what I have called “personal experience” is something like a combination of sensory perception, introspection, reason and memory. Each of these is what I call a *form of knowledge*. Each of these sources is flawed and admits of common mistakes, but the pathologies common to each individual form are unique. For instance, a disadvantage of sensory perception is that we can *misperceive* what we have experienced or *misinterpret* its meaning or importance, but an advantage is that because we have the sort of access to our own mental states described above we at least have one kind of ability to evaluate our experiences in a qualitative way by which we cannot evaluate testimony. Similarly, an advantage of reason is that, when done correctly, we can achieve a high degree of consistency between beliefs embraced through reason as well as a high degree of agreement between individuals starting from shared premises; however, a disadvantage of reason is that its utility depends on the validity of those premises, which cannot be ascertained by reason alone. Because each of these forms of knowledge is *useful* but *imperfect*, each carrying its own unique combination of advantages and disadvantages, consilience between forms is a highly desirable epistemic quality. This is another reason why personal experience is such a valuable asset when

³⁴³I have chosen to use the term “form of knowledge” rather than the term “source of knowledge,” the latter of which is preferred by epistemologists, because the term “source” begs potential confusion due to its broad utility in other passages in this chapter. Since the term “source” is used so broadly and its meaning seems more intuitive in those other passages, I have chosen to use it there in that intuitive way and fall back on the word “form” when discussing the various basic means by which we acquire knowledge, even though I acknowledge the latter term is not entirely satisfying.

it comes to evaluating testimony. Imperfect though it may be, the beliefs we form on the basis of our personal experiences are not likely to be subject to exactly the *same* sets of pathologies as beliefs adopted through testimonies, which gives beliefs that have *both* testimonial *and* experiential corroboration a form of resilience through diversity that is lacking in beliefs adopted only on the basis of strings of testimony by unknown others. This is the theory behind my saying there is *strong* consilience to the degree testimony can be corroborated by personal experience while there is *weak* consilience to the degree testimony can only be corroborated by other testimony is because personal experience, and why I think our intuitive preference for relying on the former as a basis for important decisions whenever we can get it is often justified.

Strength of Consilience: The degree to which a testimonial claim can be corroborated by personal experience, including experience with a given testifier or set of testifiers

Weakness of Consilience: The degree to which a testimonial claim can be corroborated only by testimony from unknown others

Section Two: The Epistemology of News

The Impossibility of Strong Consilience for Most of the News

Let us return to the topic of news. For a great bulk of both news and *the* news we encounter, we simply cannot get *strong* consilience. That is, most news is brought to us through chains of testimony whose reliability we cannot estimate except by consulting other testimony by testifiers whose reliability we also cannot except by consulting other testimony...and so on. In the overwhelming majority of cases, then, when it comes to telling which news is reliable and which is not, we are all in the same boat: the boat of *weak* consilience. This applies to every single individual who now exists and who has ever existed. It applies to every single reporter and editor at *The New York Times*. It applies to chief executives, legislators, members of intelligence branches, and every other member of government. It applies to the executive boards of Facebook and Google as well as anyone they might choose to employ, including “fact-checking” organizations like Snopes and Politifact. All of them experience time as a linear progression, are

bounded in physical space, and bear a perspective that is bound by those limitations just like the rest of us. None of them can be in more than one place at one time or personally witness more than the barest fraction of the events reported in the news or become personally familiar with the character of any but the tiniest fraction of the sum total of individuals reporting them. As Walter Lippmann wrote, “one common factor” characterizes the response of lay citizens and political leaders alike to the events they perceive in their time and place. That common factor is “the insertion between man and his environment of a pseudo-environment,” i.e. a simplified and flawed model of the world based on limited and imperfect information.³⁴⁴ The same limitations apply to the knowledge anyone can possibly acquire with the inner workings of the institutions dedicated to the discovery and spread of the news. Nobody can be in enough places for enough time to gain this sort of familiarity with any but the barest fraction of the internal workings of the sprawling, interconnected, bureaucratized apparatuses that together produce what we have come to call *the* news. Therefore, *all* of us, *including* every single individual who works for those very news-producing organizations, rely for the vast majority of our assessments of the reliability of the news, and the institutions reporting it (including, as may be, our own), on whatever tools we can cobble together to deal with conditions of *weak* consilience.

I stress again that my choice of the term “weak” consilience does *not* mean we should not have *any* confidence in either news or *the* news. Nor does it mean there cannot, in principle, be some people who are better than most others at judging which news is reliable and which is not. As I argued above, many of us in the 21st century use weak testimonial consilience to excellent effect in numerous decision-making areas in contemporary life, and undoubtedly some are better at using the 21st century's tools for doing so more effectively than others. But while weak consilience's weakness does not mean we should all turn into skeptics about the news, it *does* mean the confidence we attach to the beliefs we form on the basis of most news we encounter ought to be tempered in a way that appropriately takes into account the inevitable precariousness that must necessarily infuse any situation in which fallible creatures like us so heavily rely on a single form of knowledge. And, as I shall argue presently, it also undermines the idea that *anyone*, be they professional journalists, members of government, employees of Facebook or Google, or “fact-checkers,” deserves to be treated as an “expert” in the sense that citizens ought

³⁴⁴Lippmann (2012 [1920]), 10-11

generally to feel compelled to defer to their judgments about which news is reliable and which is not rather than relying on their own.

I believe most careful thinkers' approach to the news indicates that they already understand and acknowledge the epistemic precarity we are placed in when we find ourselves in conditions of weak consilience. Very few people who really give the matter some thought will opine that we ought to give any particular source of information, including even their preferred news source, our full and unqualified trust. While many people are prepared to accept the veracity of the *most* seemingly *uncontroversial* aspects of a news story – for instance, whether a battle reported on by the news *happened at all*, or whether protests *of some kind in fact occurred* downtown in response to recent police violence – the careful thinker's confidence often wanes, and ought to wane, when it comes to the *details* of that story. The careful thinker always wonders, for example, who *really* fired the first shot and whether adequate provisions were *actually* made to protect civilians, and suspects that even the version of the story provided by their favored source might be subject to bias or simple good-faith error on the matter. Similarly, the careful thinker ought always to wonder whether the protests covered by the news are more fairly characterized as “peaceful” or more appropriately be characterized as “rioting,” and whether the “extremists” the legacy news reporters' cameras fixated on were an accurate representation of the average protestor or instead a handy caricature selected by camera wielders to push a given narrative or at least gain ratings by arousing outrage and fear. Under conditions of weak consilience, our confidence in the story we eventually accept ought always to be tempered by such reservations, *not* due to a cynicism that sees everyone they don't know as corrupt or to a skepticism that sees the world as fundamentally unknowable, but instead as a simple matter of prudence that acknowledges the unavoidable precarity of depending for our beliefs on strings of testimonies disseminating from individuals and institutions whose habits, practices and character we cannot personally assess.

The same is true in the opposite direction. When we *do* have corroboration *beyond* the news reports publicly disseminated that helps us color our impression of the covered event – for example, if we ourselves or some well-known friend whose judgment and character we trust personally went downtown during the same protests reported by the news and witnessed the composition and behavior of the crowds there – we ought to consider ourselves as having a

meaningfully stronger basis for beliefs about the nature of those protests. This is not because personal experience is an infallible form of knowledge. It is obviously not. A single person's experience cannot capture the whole of even such a comparatively small and local event as a downtown protest. Nor is it because, fallible though it may be, personal experience ought always to trump testimony as an informant of our beliefs. Sometimes our personal experience on a given matter is so limited that, even though it is available, we can be justified in placing our trust in testimonial accounts that contradict our experience, reasoning that our "sample" of reality is not representative. The reason personal experience, or that of someone whose character and tendencies we know well, is so valuable in instances like the protests mentioned above, is that our personal experience, either with the situation (protest) itself or a given testifier describing the situation, gives us an additional angle by which to evaluate the situation and/or other testimonial accounts of it. Such additional angles can be extremely valuable epistemic resources, making up for much that a news report may lack. Our personal experience with a protest may, for instance, give us a sense of whether the news cameras were focused on the most representative and/or relevant aspects of the protest. Likewise, even if we could not go to the protests ourselves, our ability to use our knowledge of the judgment and character traits of friends we have known for a long time far exceeds our ability to estimate from a distance the likely judgment, character, motives and incentives behind the reports disseminated by news reporters we do not personally know, working for companies with which we have no experience, telling us stories about events we do not witness, involving people and places with which we are largely unfamiliar. We are in a better place epistemically when we can use our personal experience of a given situation, or our knowledge of our friends' judgment and character to assess the reliability of their testimony, and then compare that personal experience or that testimony to the accounts disseminated by the news, than we are when we do not have such comparative tools. Moreover, if we *frequently* had such tools available to us about the situations covered by a given source of news, we should be able to use to develop an improved sense of the *general* reliability of that news source, at least on such local matters, over time. However, in the overwhelming majority of cases and for the overwhelming majority of the news, we do not have such tools. In the overwhelming majority of cases and for the overwhelming majority of the news, then, we are unable to get more than *weak* testimonial consilience for the news, and are therefore unable to escape the essential precarity

that characterizes all beliefs that depend preponderantly on a single form of knowledge. Accordingly, we ought to temper our confidence in the beliefs we form on the basis of their reports. And, again, this applies to *everyone*. This means the limitations we place on our own confidence in the beliefs *we ourselves* derive from the news ought to temper our confidence in *anyone else's* opinion about which news reports to believe.

How *much* ought we to temper our confidence? Even if I thought a useful rubric for answering that question in any detail could be constructed, which I do not, I would not try to compose it here. What I *will* do, however, is note that a great deal of politically relevant news involves situations that have at least one of the following two qualities, and usually both. First, the situations covered by most politically relevant news are *complicated* in the sense that gaining an accurate view of all important facets of the situation – what I call the *full story* of the situation – requires a series of inferences whose reliability is very difficult, even practically impossible, to evaluate to any appreciable degree of certainty. In many cases, I believe the complexity of the situation and the unreliability of the sorts of inferences required to flesh out the *full story* of it in our heads give us good reason to suppose our confidence in any particular source's having gotten the story “right” ought to be quite limited indeed. The second relevant quality of our inferences regarding the news is that, unlike eclipses and the operations of mechanical principles, the plausibility of situations covered by most politically relevant news *cannot be reliably deduced and/or predicted by reference to established principles*, and therefore the notion of “expertise” is of limited – if indeed any – value when it comes to identifying those who are better or worse at telling which news stories and/or sources are reliable and which are not. I will discuss the implications of each of these qualities in turn.

The “Full Story” and its Importance: The Idlib Chemical Weapons Attack as an Example

I have so far often referred back to the working example of a gardener's decision about where to buy her seeds. While such a choice may excite the anxiety of those especially predisposed to fretting, for many gardeners, this decision is not exceedingly difficult. At any rate, the difficulty of that decision-making scenario is dwarfed by the difficulty of many of the decision-making scenarios faced by citizens in contemporary democracies. A few aspects of that scenario are worth emphasizing. First, (a) *the decision is relatively straightforward*. In order to

make a good decision about which seed supplier to use, most of us only need to know a few specific things related to our own priorities, such as which type and variety of plant we desire and the germination rate and average health of the plants grown from that seed stock. Second, (b) *the stakes are comparatively low*. Buying the wrong type of seed or seed that does not germinate may set us back financially and foil our attempts at creating the neighborhood's best garden, but that is a far cry from the sorts of stakes entailed in many important political decisions. Third, (c) by trying out the seeds for ourselves we are able to get *personal experiential feedback* to help us see for ourselves whether we made a good or poor purchase, and consequently whether it was wise to rely on whatever testimonial resources we may have used in this particular instance. Without this personal experiential feedback to give us strong consilience, we would have no good means of determining, over time, what degree of confidence we ought to place *in* those testimonial resources as a general rule.

In all these ways, the decision about which seed supplier to use stands in stark contrast to the epistemic task of deciding what, exactly, to believe about most politically relevant news we encounter. Take, for example, the chemical weapons incident that occurred in Idlib, Syria, on April 4th, 2017. On that date, an explosion in Khan Sheikhoun, a town near Idlib, Syria, was followed by hundreds of citizens in the vicinity struggling to breathe, frothing at the mouth, and manifesting other signs consistent with the release of chemical nerve agents. Over 80 died in the incident's immediate aftermath and video showing the suffering of victims was released on the internet shortly afterward. Immediately, Western news sources and governments accused the Assad regime of deliberately using chemical weapons to attack the citizens of Idlib. These accusations, in turn, were denied by the Assad regime and their allies, who argued instead that a bomb dropped from one of their planes had hit a warehouse where the rebels they opposed were storing chemical weapons.

How ought a 21st century democratic citizen go about deciding what to believe about events like this? Let us begin by considering the epistemic task faced by most of us in determining *whether* a chemical weapons incident occurred in Idlib *at all*, as nearly everyone who follows global political news with any degree of regularity believes with what is sure to be a high degree of confidence. I argue that their confidence in that belief is justified – or at least *benign* – for at least two reasons. First, there is an *overwhelming* amount of testimony to that end

disseminating from many different, and by all indications unconnected, sources. This means we have a *great deal* of *weak* consilience for this belief, and I have already argued that, though strong consilience is preferable, weak consilience can be and often is a legitimate grounds for doxastic confidence. Those grounds are strengthened by a second consideration, which is (a) the *straightforwardness* of the epistemic task involved in judging *whether* a chemical weapons incident occurred in a given place at a given time *at all*. The evidentiary trail left by chemical weapons attacks is not especially ambiguous, in this case involving dozens of people in the immediate vicinity of an explosion collapsing in states of near-paralysis with constricted pupils, foaming at the mouth, vomiting and choking despite there being no smoke or any other discernible toxic fumes around.³⁴⁵ It is difficult to imagine numerous witnesses being wrong about the existence of such symptoms, and the inference from those conditions in a war-torn area to a chemical weapons incident having occurred is not difficult, especially given Syria's history of involvement with chemical weapons. Finally, (b) the *political stakes* of believing *that* a chemical weapons attack occurred *at all* are not very high. By this I mean that such knowledge *by itself* does not carry any clear implications for public policy, save perhaps for sending aid to the suffering victims.³⁴⁶ Insofar as we are deciding *whether* those attacks happened *at all*, then, the *abundance of weak consilience*, the *straightforwardness* of the evidentiary trail, and the *comparatively low stakes* entailed in adopting the relevant belief that together mean that the high degree of confidence so many news readers place in the belief that the Idlib attacks happened *at all* is either justified or, at the very least, benign, and either way not worth mounting much of an objection to. And, indeed, these same features are why we should not be surprised that practically no one has denied, even in the attack's immediate aftermath, *that* a chemical weapons incident of some sort *did* in occur in Idlib, Syria, on Tuesday, April 4th, 2017.

That we can be justified in reaching conclusions like to a reasonably high degree of confidence, thus staving off epistemologists' eternal bogey-man of skepticism, is the good news. The bad news is that just determining *whether* such an incident occurred *at all* is very far from the entirety of the epistemic task that lies before citizens of Syria, the United States, or any

³⁴⁵“Syria Conflict” (2017)

³⁴⁶In truth, sadly, even this decision is made considerably more difficult due to the difficulty of ascertaining with any degree of certainty where one's money is going and how efficiently and justly it is being used, as explained in sobering detail by Leif Wenar (Wenar 2010).

concerned country in order to determine *what to do in response*. At minimum, to answer that question, the citizen must additionally determine who was *responsible* for that chemical weapons attack. It is also relevant whether they used those weapons *on purpose* or whether the chemical fallout that caused most of the damage was some freak accident. Additionally, it would be relevant to know whether that attack was fully *unprovoked*, or whether it was done in response to prior chemical warfare waged, or on the verge of being waged, by the other side. In other words, it is determining what to believe about the *full story* about what happened in Idlib that matters, not just determining *whether* a chemical weapons attack happened at all. And that epistemic task is different in nearly every meaningful respect from both our working example of deciding which computer to purchase and the example of deciding whether to believe a chemical weapons incident occurred in Idlib *at all*. Practically no one from any part of any political spectrum denies *that* a chemical weapons incident occurred in Idlib on April 4th, 2017; however, multiple sources question the claim that the Assad regime intentionally carried out that attack.³⁴⁷ In fact, all the details mentioned above as vital to the task of deciding what to do in response to the Idlib attacks were under dispute in the attacks' immediate aftermath. As mentioned above, most mainstream news sources in the West rushed to depict the attacks as an intentional use of chemical weapons by the Assad regime, but Syrian and Russian officials claimed that the chemicals actually belonged to the rebels, had been set off when the warehouse in which they were stored was struck by a stray bomb, and were now being used as a political ploy to get the international community to intervene on behalf of those rebelling against the Assad regime.³⁴⁸ These complications make determining the *full story* about the Idlib attacks a considerably more difficult epistemic task than simply deciding whether such an attack occurred *at all* for many reasons. First, weak testimonial consilience is less of an asset. Those who agree with the claim that the Assad regime is responsible tend to be officials in Western governments or reporters employed by mainstream news organizations all of which share a certain range of values and perspectives that have long made them hostile to the Assad and Putin regimes, and vice-versa. I tend to share those values and perspectives (and that hostility), but I also believe that this

³⁴⁷Obviously, these include both the Assad and Putin administrations, but the case against reflexively attributing responsibility for chemical weapons attacks to Assad forces has also been made by investigative journalist Sy Hersh (Hersh 2013), as well as Democratic member of the House of Representatives Tulsi Gabbard and Republican House member Thomas Massie (Mills 2017).

³⁴⁸Barnard and Gordon (2017)

uniformity of perspective between sources limits the value of agreement *between multiple* sources, as because any flaw endemic to that shared perspective may easily undermine all of these sources' interpretation of events in Idlib. Second, unlike the earlier example of a gardener deciding which seed supplier to use, the stakes are *high*. Whether or not the West intervenes militarily in Syria has life-or-death implications, and public pressure to intervene is likely to be a non-negligible influence on elected leaders. Citizens' beliefs about who is responsible for the Idlib attacks and to what degree are an important factor in determining how much pressure is placed on Western governments to intervene, so the stakes are very much life-or-death.³⁴⁹ Third, and finally, for anyone outside the respective militaries, the citizens in the vicinity of the warehouse supposedly storing the chemical weapons, and perhaps scientific specialists familiar with the properties of the sarin gas thought to be involved in that attack, the epistemic task entailed in deciding who is responsible and to what degree is *complicated* because it requires citizens to make a *series of relatively speculative inferences*. Inferring that a chemical weapons incident of some kind has occurred upon seeing dilated pupils and foaming mouths in a war-torn city is relatively straightforward. Inferring the exact origin and intended use of those chemicals is, for most of us, considerably less so. For the majority of us who do not have direct familiarity with the inner workings of either military or with the warehouse in question, we are forced to base our views in the latter case on a series of inferences drawn from our working notions of the past behavior and likely motives of both sides, our antecedent views of the morals and material capabilities of the actors involved, our estimations of the trustworthiness of our various sources of information such as news agencies, UN commissions, and so forth. These resources are not worthless and I suspect many of us are able to use them to form pretty reliable beliefs an appreciable amount of the time. However, the force with which this sort of evidence *demand*s a given inference is relatively weak compared to the force with which the inference that a chemical weapons attack occurred *at all* is demanded the evidentiary trail described above or, to return to our earlier example, the evidentiary trail left by the performance of our chosen laptop.

³⁴⁹Indeed, such considerations may well have played a decisive role in the Trump administration's decision to sanction the bombing of the Al Shayrat airfield, where at least six people died (Said-Moorhouse and Tilotta 2017).

I believe all of these factors ought to temper our confidence in any particular story about the chemical weapons attacks in Idlib, Syria, on April 4th, 2017, especially in the immediate aftermath during which the “fog of war” tends to be thickest. They ought not prevent us *from* forming an actionable opinion on the matter, but they certainly recommend more restraint than was commonly shown in the discourse in the attacks' immediate aftermath, where many commentators who had no personal knowledge of any of the decision-makers involved or the on-the-ground conditions treated the proposition that the chemical weapons attacks were intentionally perpetrated by the Assad regime as self-evident and piled public pressure on the Trump administration to immediately respond, which it eventually did by bombing the Al Shayrat airfield.

Defense Against Charges of Excessive Skepticism

I suspect many readers will feel a strong urge to resist my position here, thinking it excessively skeptical. However, I ask such readers to subject that urge itself to scrutiny and consider whether it is a sign of my having gone astray or, instead, of a sign of a sort of generalized miscalibration in our epistemic approach to the news. I believe what I have argued here is consistent with most of our views about the proper epistemic approach to decision-making even on issues with considerably lower stakes in other realms. When someone is accused of even a comparably minor crime like theft, we demand they be given a chance to defend themselves in a trial governed by numerous procedures designed to improve the epistemic circumstances faced by those asked to issue judgment. Jurors or judges, for example, have a chance to personally experience, over an extended period of time, the testimony of such eyewitnesses as can be found, listen to rigorous arguments formed by representatives of the accusing and defending parties, and examine evidence deemed both relevant and admissible according to well-defined standards. Only after such a careful and rigorous process has been undertaken do we believe *sufficient confidence to be warranted* in the conclusions derived from that process to *justify* rendering a decisive verdict, even when the consequences of that verdict are as little as a monetary fine or a short stint in prison, which, while not negligible, cannot rival the normative stakes of decisions that bear implications for the life and livelihood of several, or even thousands, of people. Indeed, even in many of those cases that undergo such rigorous

procedures as those entailed in trial by jury, we *still* suppose the evidence so scrupulously gathered and the arguments so carefully marshaled to frequently give insufficient epistemic grounds to justify confidence in a verdict one way or another, and so have formalized a “default” judgment of “innocent until proven guilty” to clarify what to do in those common cases where that process is deemed insufficient. Like many others, I believe this approach to criminal procedure is absolutely indispensable. It is justified by the *stakes* of the decisions entailed in criminal accusations as well as the *epistemic conditions* required for human decision-makers to have any hope of reaching conclusions in a way that justifies the level of confidence we ought to demand of decisions about whether or not to convict someone of a crime. Most of us would find it unconscionable if, instead of going through all this rigmarole, we simply had judges or jurors read a few news reports or listen to the public statements of government officials, none of whom they knew or had the chance to become familiar with in any way, and many of whom may well have a conflict of interest in the matter at hand. If we would be appalled at a legal system that treated even comparatively minor criminal accusations in this way on the grounds that the stakes demand a higher epistemic standard, why should we so easily accept legacy news reports as sufficient grounds to justify the public's putting pressure on politicians to adopt public policies likely to entail far greater consequences, such as launching planes and bombs in Syria?

Now, obviously, we cannot always meet the same epistemic standards attained via trial-by-jury with regard to all the beliefs on which we must render judgment as citizens of 21st century democratic societies. I have already acknowledged this in the previous chapter in the sub-section on *tragedy*. But I also pointed out in that sub-section that making oneself *too readily* complicit in tragedy is just as much a moral failure as an unwillingness to admit that the conditions of tragedy often necessitate acting even under less than ideal epistemic circumstances. I believe one way we commonly demonstrate excessive readiness to make ourselves complicit in tragedy in this sense, and one encouraged by Progressives, is by lending too casually and without much scrutiny a high degree of confidence in the stories of the world brought to us through the news. This easy lending of confidence is at odds with the more demanding epistemic standards we deem – in my view correctly – requisite to meet the epistemic standards we demand in order to justify decisions of even considerably lower stakes made in other arenas. If the reader's impulse, then, is to scoff at my chosen example of determining who was responsible, and

in what ways, for the Idlib chemical weapons incident, on the grounds that that example demonstrates excessive skepticism on my part, I respectfully suggest readers consider whether such scoffing might be emblematic of precisely the problem I am identifying. I ask such readers to reflect on why they demand such high epistemic standards for decision-making in comparatively low-stakes circumstances but settle for so much less when it comes to beliefs that inform people's public policy preferences, where the stakes are far higher.

The example of the chemical attacks in Idlib is, in my view, far from an outlier among politically relevant news stories in all the respects I have emphasized. That is, for the vast majority of the news we encounter, forming an accurate view of the *full story* entails (a) relying entirely on weak consilience, or testimony by testifiers whose general reliability we cannot assess by recourse to any *direct experiential feedback*, (b) *high stakes* in terms of the public policy implications, and (c) making a series of comparatively *complicated and speculative inferences* on the part of the citizen. Far from being unique to the Idlib chemical weapons attack, in all these respects I believe the Idlib example to be quite representative, if anything perhaps even comparatively simple and straightforward. A cursory scan of just the headlines on the front page of today's *Washington Post* reveals it to include stories that invite citizens to make judgments on such matters as whether the Trump administration is jeopardizing American democracy by refusing to cede electoral victory to Joe Biden, the degree to which America's hospital system is being overloaded by the COVID-19 pandemic, and the scope of the culpability of Catholic church leaders, including Pope John Paul II, in enabling the sexual abuse scandals that have rocked the church for the last few years.³⁵⁰ Most of us have no more resources for forming an accurate picture of the *full story* in any of these instances than they do regarding the chemical weapons incident in Idlib, and the same is surely true of most of the stories beyond the front page. This does not mean we should not *form* beliefs on these matters and endorse public policies on their basis, but it *does* mean our confidence in our own conclusions – and those of anyone else – ought to be limited. This will have important implications for how we treat those who come to different conclusions than our own, which I shall explore later in the dissertation.

350The *Washington Post*, November 11, 2020, front page

The Absence of Identifiable Expertise about the News and its Implications

The final aspect of our epistemic relationship to the news I will dwell on in this section is what I call the *absence of identifiable expertise in the epistemic domain of news*. For my purposes, the term “expert” will be taken to mean the following:

Expert: An individual or institution whose understanding of a given epistemic domain is sufficiently greater than that of the lay person to justify routine deference by lay persons to the claims of that individual or institution on matters falling under that domain (so long as that individual or institution is thought likely to be trustworthy in the situation at hand).

I have chosen this definition, which ties the notion of expertise to the notion of routine deference, because a central complaint of Progressives is that too many citizens obstinately refuse to accept the claims of experts on matters of political importance, such as climate change, Brexit, and vaccination. Accordingly, the notion of expertise invoked by Progressives, and democratic citizens' proper relationship to it, implies that the attitude citizens ought to take toward experts is one of deference, that is, a willingness to trust those experts' claims. It is important to note two things about this position. First, *deference* implies *difference*. The epistemic utility of relying on the claims of experts stems entirely from situations where lay citizens' own judgment differs from that of the experts, otherwise such deference would be entirely superfluous. I will explain why I think this has important implications for citizenship in liberal democracies a bit further down. Second, I take it as obvious that common-sense worries about conflicts of interest and potential abuses of power to apply in this situation as readily as they do in all situations involving deference to authority. That is, *epistemic authority may as easily be abused, and invites abuse just as readily, as other forms of authority*. It should be obvious that we cannot count on experts always necessarily having the public's best interest in mind. Accordingly, we have to watch out for situations where a conflict of interest might lead them to abuse any excessive readiness to defer to their expertise. This is why I have consistently emphasized the requirement that a testifier needs to be both credible *and trustworthy* in order to justify citizens' decision to place trust in their testimony. This means that simply *knowing better*

than lay citizens is not sufficient grounds for arguing that citizens ought routinely to defer to a given expert. To this we must add the requirement that the expert is likely to be *trustworthy* in the situation at hand.

The notion of expertise as I have defined it above is implicit in a great deal of contemporary commentary about the current and proper roles of *the* news and *fake* news in our era. In at least three ways, the arguments of Progressives strongly imply that lay citizens ought to consider the epistemic domain of *the* news to be a domain in which the concept of expertise applies. First, they explicitly *call* journalists “experts,” usually in contexts where it is clear they mean that label to justify the sort of attitude of routine deference I describe above.³⁵¹ Second, Progressives strongly imply democracy would be better off if lay citizens adopted a routine attitude of trust toward the legacy news media and refrained from making decisions about which sources to trust based on their own judgment. This is evident in the frequency with which Progressives cite decades-long declines in trust in the mainstream news media as a sign of democratic decay³⁵² as well as the regularity with which they lambast lay citizens who *do* use their own judgment whenever that judgment leads them to place their trust some online source *other* than the legacy news media.³⁵³ Third, many Progressives endorse the censorship of fake news by some set of authorities, be they governmental authorities³⁵⁴ or private employees of platforms like Facebook and Twitter.³⁵⁵ It is difficult to imagine a defense for that position if its advocates did not believe in the existence of some identifiable group of people who possessed judgment about which news “really” deserves to be trusted that is sufficiently superior to that of the lay public to justify handing control over such decisions to them rather than leaving them to the public itself.

351See, e.g., Goldman (2010), Nichols (2017, ch. 5), Persily (2017)

352E.g. Goldman (2010), Schiffrin (2017), Deb, Donahue, and Glaisyer (2017)

353E.g. Dahlgren (2018)

354Censorship of fake news by governmental authorities appears to be supported, for instance, by communications scholar Kathleen Hall Jamieson, who praises French President Emmanuel Macron for threatening members of the French press with legal punishment if they published “hacked content” about Macron in the days preceding the 2017 election (Jamieson 2018, 11-12).

355Censorship of fake news by platforms such as Facebook and Twitter enjoys widespread support, particularly on the political left, where House Speaker Nancy Pelosi and Senator Elizabeth Warren have repeatedly berated Facebook in particular for failing to do so. For an example of support for this position in the academic world, see Levinson (2019).

I believe Progressives go wrong both in applying the notion of expertise, as defined above, to the epistemic domain of news³⁵⁶ and in arguing or implying that lay citizens ought to routinely defer to any set of epistemic authorities' claims about which news is "true" and which "false." Let us take these arguments one by one, begin with the argument that the notion of "expertise" ought not apply to the epistemic domain of news. Some epistemic domains are governed by rules and regularities that are unknown to the lay public but are comparatively well-understood by specialists in that domain. In many cases, this explains why people who are very familiar with a given subject area are so much better at vetting testimony than people who are not. For example, a quantum physicist is much better positioned to verify or contradict the claims my brilliant but eccentric stoner friend Andy occasionally makes about the way matter works at a fundamental level than I am. This is because, although quantum physics is far from a perfectly understood science, it *is* governed by at least *some* rules that are well understood by quantum physicists, and a quantum physicist has much more thorough knowledge of those rules than I have. The only "rules" about quantum physics I can use to vet testimonial claims in that epistemic domain are that the behavior of quantum particles is *probabilistic* rather than *deterministic* (I think) and that quantum physics is *weird* and therefore I oughtn't too readily dismiss any particular claim Andy makes about the behavior of quantum particles no matter what he's been imbibing lately. While not utterly useless, these rules of thumb I use for navigating the epistemic domain of quantum physics are very far from qualifying me to make reliable judgments about whether Andy's claims about the behavior of sub-atomic particles are likely to be true. For a trained quantum physicist, the situation is different. All they have to do is consult their existing stock of knowledge about the rules that reliably govern the behavior of quantum particles and see whether Andy's claims stack up with them. Because quantum physicists are aware of certain well-established and highly reliable principles about the way matter works, they can use this knowledge to assess, with a high degree of justified confidence, whether Andy's claims are likely to be true. And because I believe quantum physics is an epistemic domain governed by such rules, and quantum physicists to have a better grasp of them than me, it seems

³⁵⁶The "as defined above" part here is important. I do not deny that journalists' training helps them develop skills that, when employed properly, are useful in the pursuit and dissemination of knowledge. I *do* deny, however, that this sort of training justifies lay citizens deferring to the opinion of journalists about what news ought to count as "reliable" in the same way an engineer's training justifies lay citizens deferring to them about how to properly build bridges.

reasonable for me to be ready to defer to their opinion on matters of quantum mechanics – so long, at least, as I also believe the quantum physicists I am consulting to also be *trustworthy* given the situation at hand, and not to be influenced by some conflict of interests, such as the desire for additional research funding.

The situation regarding the epistemic domain of *news* could scarcely be more different. “News” can be about anything of any interest whatsoever to any appreciable number of people, which effectively means the epistemic domain of news encompasses practically every subject matter known to humanity, from quantum physics to politics to economics to croquet. It is extremely doubtful that any particular individual or institution understands whatever “rules” govern that vast arena of affairs sufficiently better than lay citizens justify citizens' being willing to routinely abandon their own judgments in deference to the former's opinions on all the matters covered in that sweeping epistemic domain *even if* citizens could somehow be assured of that individual or institution's good intentions. Beyond this, as I shall argue below, even if such experts existed *as a matter of fact*, the fact *of* their expertise could not be demonstrated publicly to any satisfactory degree. Though physicists undoubtedly still have much to learn about quantum phenomena, the degree to which their understanding of the rules that govern quantum mechanics outstrips that of the lay person, allowing them to determine with a much higher degree of accuracy what may or may not be true in the epistemic domain of quantum physics, is far more evident than the superiority of any identifiable individual or group's ability, compared to that of the lay citizen, to use their understanding of whatever rules may govern the vast array of human affairs covered by the news to say which news *may* or *may not* be true about *the world*.

I will return to this point in greater depth later, but for now I just want to point out one implication of the last few sub-sections. Together, I take the observation that we can get only *weak* testimonial consilience for the overwhelming majority of news we encounter, together with the general *meagerness of our epistemic resources for evaluating the news* and the *absence of identifiable expertise in the domain of news* to constitute a *prima facie* case for substantial epistemic humility when it comes to evaluations of the scattered collection of testimonies that together comprise the news. This does not mean we should refuse to form beliefs on the basis of the news or even refuse to commit to substantively important and consequential political policy positions on its basis. However, it *does* mean that we ought to recognize that for *all* of us, *most*

of the time these beliefs and policy positions are adopted under epistemic circumstances that are very far indeed from ideal. Accordingly, in most cases our conviction in them, our readiness to chastise and demean others whose opinions differ from ours, and certainly our willingness to embrace public policies that hinge crucially on our conviction that *our* evaluations are reliable while those of other fellow citizens who *disagree* with us are not, ought by default to be tempered by an acknowledgment of the significant and pervasive uncertainties entailed in the formation of reliable beliefs under such circumstances – a layer of uncertainty beyond the chronic epistemic fallibility that characterizes the human quest for knowledge generally.

One crucial implication of this line of reasoning is that we have *prima facie* epistemic reasons to be very suspicious of *any* individual or institution that claims to have a special ability to tell which news is “true/real” and which is “false/fake.” For the same reasons, the argument so far already constitutes a strong presumption, again on epistemic grounds, against the idea that anyone deserves to be treated as an *authority*, in the epistemic or (especially) legal sense, about which news constitutes “true/real” and which news is “false/fake.” I believe this case only becomes stronger once we consider in a little more detail the specific tools citizens have available to them for deciding whether and how far to believe testimony such as the news, which is a goal of the next section.

Another Kind of “Expertise”

Some might suspect that keeping the conversation at the abstract level I have so far employed might lead us to conclusions that wouldn't hold up were we to subject them to the “sniff test” of actual experience. In this section, I will try to assuage these worries by describing the three best tools I can think of that citizens may employ as a practical means of dealing with situations characterized by epistemic distance and being limited to *weak* testimonial consilience, using the news as a running example. These tools are (1) consultation of a testifier's *reputation*, (2) comparison of testimony against our *worldview*, and (3) the pursuit of *perspectival diversity between testifiers*. I doubt this is anything like a complete list of all the tools that might be employed to maximize the reliability of one's beliefs under conditions of *weak* testimonial consilience, but I do think these three are among the most important and most useful. Examination of these tools will show that (a) even though we cannot usually get *strong*

testimonial consilience with regard to those sources and reports, we as individual citizens are far from helpless when it comes to deciphering which news reports and news sources deserve to be trusted. Still, (b) the utility of those tools under *anyone's* use is *limited* and therefore, consistent with my argument above about the precarity of relying on a single form of knowledge as one's entire grounds for belief in a given proposition, the reliability of *anyone's* conclusions based solely upon use of such limited tools must be treated as being somewhat *precarious*. Therefore, (c) good use of those tools depends critically on *judgments* in epistemic domains where *no convincing case can be made that any identifiable group of persons holds a level of expertise of a sort that ought to compel us to take their claims more seriously than lay citizens'*. The tools I bring up in this section are among the best epistemic tools available to *anyone* assessing the reliability of the vast majority of the news. Nobody employed by Facebook, Google, any government, or anyone else has, in the vast majority of cases, any better tools available to them than these, and they are the same tools available to the lay citizen. Accordingly, while lay citizens ought to exercise epistemic modesty and admit that their own conclusions about which testimonies and testifiers are reliable may easily be wrong, they are perfectly entitled to feel the same way about the conclusions reached by everyone else, too, including “fact-checkers.”

My discussion here is meant both to support the line of argument prosecuted in the previous section, which denied that there were identifiable experts about the news, and to defend the case against the existence of identifiable expertise about the news in another sense. My previous case against the existence of identifiable expertise about the news was based on the argument that there is no good reason to suppose that any individual or institution understands whatever rules may govern the vast epistemic domain of news sufficiently better than that of lay citizens' to justify lay citizens' routine deference to them, and that even if some such individual or institution *did* exist, its expertise could not be demonstrated to the public with sufficient clarity to justify its being treated as a *legitimate authority* by the public. Some may have acknowledged this point but still insisted that the notion of expertise applied to the epistemic domain of news, not because “news experts” understood the *rules* that *governed* the news, but instead because they are *more practiced* or *more skilled at* using the tools we have available to us for assessing the reliability of the news. I believe this is line of argument, too, is incorrect. For each epistemic tool I will mention, I will argue that although it *can* be a useful means for vetting

testimony, its utility is limited and its conclusions precarious *no matter who uses it*. I will also argue that, beyond this basic precarity, we cannot reliably identify whatever individuals may be more skilled at employing these tools. These considerations give lay citizens even more reason to resist expectations of routine deference to would-be “authorities” about what news they ought to consider reliable, and certainly to the imposition of censorship by anyone claiming to possess this kind of expertise. Indeed, I will argue that an examination of each of these epistemic tools reveals the need to *encourage* in lay citizens to *assertively scrutinize* testimonies, testifiers *and* those presented to them as would-be epistemic authorities *in the light of their own judgment* in a way that is actively discouraged and disparaged by Progressives, who cite such assertiveness on the part of the citizenry as a pathology characteristic of the post-truth movement.

Section Three: Reputation

One tool we may use to assess the reliability of testimony for which we can obtain only *weak* testimonial consilience is attention to an individual or institution's *reputation for credibility and trustworthiness in a given epistemic domain*. In this sub-section, I will first argue that the idea that reputation is a good reason we ought to trust the legacy news media either (a) undermines other arguments commonly espoused by Progressives or (b) fails to resolve the problem it seeks to address. I will then argue that, irrespective of the arguments of Progressives, while consultation of reputation *can* be an *acceptable* way of vetting testimony in the absence of anything better, beliefs formed on its basis are still *precarious* in the way I have argued is characteristic of beliefs formed under conditions of weak consilience. This ought to limit the confidence we are willing to place in them and, accordingly, our readiness to castigate those who reach different judgments than our own about whether an individual or institution's reputation merits our trust.

Having established a reputation for trustworthiness is frequently cited as a reason we ought to trust the legacy news media – that is, by the rare few who feel the need to defend that position at all.³⁵⁷ It is explicitly appealed to, for instance, by James Carson of the *Telegraph's*

³⁵⁷ Much more common is for the opinion that we ought to trust the legacy news media simply to be stated without explanation, which is a tendency among those I take to be my opponents which I shall have occasion to point out on a number of occasions and on a number of different topics in this dissertation.

narrative of fake news' emergence in the past decade. Carson argues that in the past, building a “large audience took much longer, and because it was expensive to acquire and built on trust of information, publishing fake news would be damaging to reputation and thus have economic consequences.” However, in recent years the “costs of publishing (via WordPress) and distributing (via social networks) approached zero,” which makes reputations “far more expendable.”³⁵⁸ The implication of Carson's argument would seem to be that we have good *prima facie* reasons to place our trust in the claims of legacy news organizations over practically *any* informational source that emerged in the digital era and operates primarily online, a sentiment that fits well with the attitude of friendliness toward legacy news media – and long-established institutions in general – and suspicion of newly emerged online alternatives often expressed by Progressives.

Setting aside the objection that it is not immediately clear *why* lowered costs of publication should be expected to drastically diminish online publishers' desire to establish a good reputation,³⁵⁹ I believe Carson's argument that we ought to trust legacy news because it enjoys a good reputation sets up a dilemma for Progressives. Carson's argument seems to be based on *incentives*. Since their success as a business relies on their continuing to be seen as a source of accurate and germane information, so the argument seems to go, we should the legacy news media to do whatever they can to make sure what they publish is accurate. I find this to be a common defense of trust in the legacy news media in casual conversation, especially among those most worried about the effects of 21st century communications technologies on democracy, though the propensity for Progressives to simply *assume that* the legacy news media *just is* (self-evidently?) more reliable than online-era alternatives is so dominant³⁶⁰ that it is rare to see it articulated explicitly in print. Common as it is, there is an assumption hidden in this line of reasoning that is worth pointing out. That assumption is that news consumers will *reward accuracy* by improving the esteem in which they view news agencies that give accurate reports and *punish inaccuracy* by lowering their esteem in which they view those that do not. In other words, the idea that we ought to trust news agencies that enjoy a reputation of reliability

358Carson and Cogley (2017), their parantheticals

359 After all, the aim of gaining *consistent web traffic* would appear to be common to many websites regardless of input cost. If this is what incentivizes the legacy news to tell the truth, why should we not expect it to incentivize other online entities to do the same?

360Deb, Donahue, and Glaisyer (2017), Schiffirin (2017), House of Commons Report (2018), UNESCO (2018)

implies a relatively optimistic opinion of the *judgment of the public* when it comes to assessing the reliability or unreliability of news sources. If we are to consider having an *established reputation* for reliability as a good reason to trust the legacy news media, as Carson seems to be arguing, then the public would appear to *already be* pretty good at judging which news ought to be considered reliable.

This poses a problem for Progressives because, as we have seen, the idea that the public is *already* pretty good at judging which news deserves to be considered reliable is sharply at odds what most Progressives appear to believe. It is the opposite of what is implied by the idea that misinformation and fake news are such big problems that they pose a threat to democracy, as well as being the very opposite of what is implied when people like Carson expend considerable time and effort publish guides informing the public about how to tell real from fake news – time and effort which would seem to be ill-used if the public already possesses sufficiently sound judgment to reliably accord a good reputation to reliable sources of information and view the alternatives with scorn. In fact, the belief that the public is *not* good at differentiating reliable from unreliable sources seems to be the very *reason* Heirs of Progressivism often view the legacy news media and other long-established institutions as indispensable to democracy and to treat free-flowing informational environments like those provided by social media platforms a threat to it. How else can we explain the view espoused by Anya Schiffrin, head of Columbia School of Public Affairs' department of Technology, Media and Communications, that disinformation spread on social media is sufficiently prolific as to “undermine the functioning of democracy globally,”³⁶¹ largely because citizens have abandoned their trust in long-established institutions like those of “education, science, and media,” which “traditionally kept 'false facts' and demagoguery at bay” in favor of online alternatives with inferior “standards of truth?”³⁶² That view seems necessarily to paint the public as not just *poor* but *catastrophically unreliable* judges of what testimony should or should not be relied upon. Why should we suppose that, given the wide-ranging exposure to *all* sorts of news, both reliable and unreliable, that is available online, the public would embrace the latter and reject the former so consistently as to “undermine the functioning of democracy globally?” This combination of views so common

³⁶¹Schiffrin (2017, 118)

³⁶²Ibid, 123

among Progressives – that the online informational environment characterized by lowered barriers to publication is a democracy-undermining cesspool of disinformation while established institutions like the legacy news media are needed to protect citizens from “false facts” – would appear to be nothing less than a profound rebuke of the public's judgment. But if the public's judgment is this poor, how can an institution's *reputation* be a good reason for us to place our trust in it? A thing's reputation is the *public's opinion of it*. The idea that reputation is a good reason to trust in an institution would seem necessarily to imply that the public is *already* a good judge of which institutions deserve its trust. This forces a dilemma for Progressives who cite reputation as a good reason to trust legacy news media. Apparently they must either continue to believe reputation is a good indicator of trustworthiness, and in so doing admit that the public is a decent judge of such, or they must double down on their pessimism about the public's judgment and find some other reason to justify trust in the legacy news media.

Given this forced choice, which will Progressives choose? I believe a great many will be likely to take the latter option. I am inclined to believe this in part because we have been working under a false assumption – the assumption that the legacy news media *does in fact* enjoy a reputation among the public for being a credible and trustworthy source of information, as Carson seems to imply. But it is not at all obvious that this is the case. On the contrary, for the majority of the last half-century the lay public's attitude toward the legacy news media, at least in the United States, whose political situation is the focus of much of Progressives' work, has increasingly become one of skepticism, wariness, and distrust, a trend that is frequently cited by Progressives as a bad sign for democracy.³⁶³ American citizens expressing a “great deal” or “a lot” of confidence in the mainstream news media hovered between 30-40% for decades prior to the election of Donald Trump,³⁶⁴ whose combative stance toward the media appears to have briefly inspired a rally of pro-media sentiment among Democrats that lifted aggregate confidence measures during the first two years of his presidency,³⁶⁵ though it is hard to say how long the latter trend will last once his term ends in early 2021.³⁶⁶ I believe that, when forced to choose between (1) agreeing with Carson that reputation *is* a good sign of reliability, and by implication

363E.g. Goldman (2010), Persily (2017), Schiffrin (2017)

364Dugan (2014)

365PEW (2019)

366Early signs suggest it may be short-lived indeed, as Americans' expressed trust in the mass media began to dip again in 2019 after receiving a small boost in 2017-18 (Jones 2018).

admitting the public is *already* a reasonably good judge of what news sources are and are not reliable, or (2) denying that reputation is a particularly good sign of reliability and citing the public's distrust of legacy news as an example, many Progressives will choose the latter.

Of course, another way of resolving the contradiction between the idea that reputation is a good sign of reliability and pessimism about the public's judgment would be to argue that reputation *is* a good indicator of a news source's reliability, but only the reputation that source enjoys *in certain circles*. But this begs the question of how to figure out *which* circles are the ones whose opinions about news sources we ought to consult. Progressives' propensity to simply *assert* which sources they prefer and to condemn the “post-truth” tendency on the part of citizens to refuse to do so seems to imply either that the reliability of *their* favored news sources ought to be self-evident or that *their own* opinions ought to be authoritative on the matter. I deny that the superiority of legacy news over online-era alternatives is self-evident. So exceedingly few things are really “self-evident” that the use of that phrase is nearly always – and is indeed in this case – an attempt to evade the need to give reasons for one's opinion rather than a reason of its own. The idea that it could ever be “self-evident” which testimonial sources are telling one the truth in such a vast and complicated epistemic domain as that covered by the news, when one can nearly never examine those testifiers' claims oneself, is even more preposterous than the already-preposterous uses to which that term is often put. That leaves us with the implication that lay citizens ought to simply *suppose* Progressives' own opinions about which testimonial sources are reliable to be authoritative, and to justify those sources' being treated with the attitude of routine deference so frequently encouraged by Progressives. I deny that we ought to do so. The expectation that we ought to simply *suppose* Progressives to be better judges of which testifiers are credible and trustworthy, and therefore to defer to those sources as a matter of course, is out of keeping with the *strong preference for doxastic self-rule* I have argued to be integral to any convincing account of normative epistemology, as well as with any liberal democracy worth defending. As I argued in chapter one, in such a liberal democracy, deference to authority, be it legal or epistemic, ought not to be expected as a matter of course, or by definitional fiat, but instead needs to be *justified through reason* to the person from whom deference is expected. The argument that its enjoyment of a reputation for trustworthiness *in certain circles*, then, ought not

to be considered a very persuasive argument for why the public ought to consider legacy news media to be generally more reliable than alternative sources that have arisen in the internet era.

So far I have objected to Progressives invocation of reputation as a reason lay citizens ought to trust the legacy news media more than online-era alternatives on grounds of *internal inconsistency*, arguing that *either* enjoying a reputation of reliability is not an especially strong grounds for trust in its testimony *or* that it *is* a strong grounds and that therefore Progressives' deep pessimism about the public's judgment is unjustified. Let us now consider the question on its own terms. How reliable *is* reputation as a guide for deciding which testimonies and testifiers to believe? I think it is obvious that it is not entirely useless. It is in part on the basis of their reputations for reliability, for instance, that consumer reports magazines and similar buying guides gain loyalty. As I have said before, many of us have used such resources plenty of times and often been happy afterward with the decisions we made on their basis.³⁶⁷ I believe this utility derives in part from the same general acuity of judgment and propensity for honesty I argued to be characteristic of most citizens' everyday decisions and disseminations of testimony, which I have cited in previous chapters as grounds for my *modest optimism* about lay citizens both with regard to their ability to navigate the online informational environment and with regard to their fitness for liberal democracy. While I believe this signifies that relying on reputation to inform one's decisions and beliefs has *some* utility, however, it must be admitted that it has its limits. It should be uncontroversial to say that sometimes individuals and institutions develop reputations they do not deserve. Sometimes, this happens because the individual or institution acquired that reputation in a former time when it was actually deserved, but has since changed its behavior and societal notions have not yet caught up. Sometimes, that reputation was not founded on good grounds in the first place, perhaps having been the result of persistent PR campaigns – i.e. chains of carefully cultivated testimony – specifically designed by the institution itself designed to cultivate such a reputation. Moreover, even institutions that *do*

³⁶⁷Of course, the situation is more complicated than this. There is, for instance, also the “priming” effect, now well established in psychology, which suggests that part of what might be happening in the case of consumer reports “successfully” leading consumers like me to “good” choices is that consumers primed by consumer reports are more inclined to view their purchase in a positive light. Additionally, there is the desire to rationalize our past decisions, which appears to be the motive behind the otherwise puzzling tendency for consumers to seek out advertisements for a given product *after* having already purchased it.

deserve a *general* reputation for trustworthiness fail to live up to it in some instances, either as an honest mistake or for more pernicious reasons.

It is important to be clear that these dangers most certainly apply to the legacy news media. There have been numerous occasions on which these testimonial sources have clearly failed to live up to their reputation among Progressives as “communicators who work in the service of truth”³⁶⁸ who protect democracy by keeping “false facts at bay.”³⁶⁹ In numerous instances, including but not limited to the Gulf of Tonkin incident and the claim that the regime of Saddam Hussein was producing weapons of mass destruction, the American mass media has acted with a seemingly unanimous lack of interest in contesting claims on whose truth hinged the justification for starting or escalating military conflicts which would claim many lives and do enormous damage to many others.³⁷⁰ Instead, in both cases, they merely absorbed and repeated the claims of authorities in the U.S. government, even though they clearly had reason to suspect those authorities had either a conflict of interest or at least biased judgment when it came to making the decision to enter or escalate the conflict. Such incidents reveal not only that legacy news media cannot *always* be relied upon to report the *full story* in the ways important to politics, but also that reputational incentives may have a dark side, too. While it may be true that the desire to maintain a reputation for accuracy and fairness is *one* set of incentives faced by the legacy news media, it is not the *only* incentive to which they must respond in order keep producing news (and profits). One of these sets of incentives is that they tell a good, and *timely*, story.³⁷¹ Another is that they maintain a reputation for being on the side of “the good guys,” where the “good guys” means the side of whoever they see their viewership as belonging to. When citizens are rallying around the flag, for example, it can be a very poor time, pragmatically speaking, for news organizations to start publishing stories illustrating the virtues of patience, emphasizing the contingency of all human knowledge, and calling into question the moral conviction that subsidizes the march to war. These efforts at holding to a higher code of honesty

368UNESCO (2018), 9

369Schiffrin (2017), 123

370For a thorough examination of this prolific tendency of the media all through its supposed “golden age” between 1945-1991, see Zaller and Chiu (1996). For one of many accounts of the many scathing critiques of the news media's approach to the Iraq War and its acceptance of administration claims about weapons of mass destruction, see Moeller (2004).

371Kathleen Hall Jamieson and Paul Waldman illustrate the way this imperative for news organizations to provide a narrative often competes with their sanction to report the truth (Jamieson and Waldman 2003).

and truth-seeking may, at times, as easily *diminish* the news organization's reputation in the eyes of its potential consumers as publishing false or misleading reports.

Nor is wartime the only occasion for failures by the legacy news media, as many other instances of oversight, negligence, or omission could be recounted, sometimes at the express request of governmental leadership whose interest is directly at stake.³⁷² For we are very far from exhausting the list of potentially perverse incentives faced by legacy news organizations. One additional incentive worth noting is that of *profit*. I doubt any reader will be shocked at the suggestion that profit motive may at times be at odds with full forthrightness, and for the legacy news media, it certainly has been. The quest for profit was why local TV stations often outsourced their “news reporting” airtime to private companies, who designed advertisements disguised to look like “news reports,” for the better part of three decades. From at least the mid-80s to the mid-00s, companies produced “video news releases,” or VNRs, in which they would create stories that contained their products and have them aired by local news broadcasters in an effort to pass their advertisements as “news.” Done up in the style of “objective” news reports and bearing the insignia of the local syndicate on which such reports were broadcast, VNRs were “virtually indistinguishable from honest news, featuring interviews with (paid) experts and voiceovers by (fake) reporters who subtly pitch products during their narratives.”³⁷³ In other words, legacy news companies during this period *themselves* regularly broadcast fake news. These were not exclusively manufactured by *private* companies, however. The U.S. government produced and covertly placed VNRs on news programs during this era, too, as the George W. Bush administration was eventually discovered to have done in a story that scandalized the public, and finally made some citizens aware that some amount of their “local news” content was actually paid advertising, when finally revealed.³⁷⁴ Still, however, despite its evident importance in terms of the implications for society's relationship to the news and its role in “protecting democracy” by keeping it “informed,” these revelations about VNRs have never gained

³⁷²The Cold War was the era during which and justification for many of these instances, but one need not reach so far back for examples. American abuse of Iraqi prisoners at sites such as Abu Ghraib were known of by legacy news media institutions for at least a year before their publication was forced, not by competitors in the news media, but by the release of photos of such abuse into the blogosphere.

³⁷³Manjoo (2008)

³⁷⁴Pear (2005)

widespread public knowledge, probably because they have never gained the attention of the legacy news media itself. It is not hard to guess why.

The fact that legacy news organizations could get away with pitching products and partisan governmental policy proposals under the guise of reporting “the news” for over 25 years³⁷⁵ with few members of the public being made aware of the practice, even to this day, must surely call into question just how much we should feel free to rely on the legacy news media's reputation as a trustworthy source of information even if, despite consistent polling results to the contrary, it could in fact be said *to* enjoy such a reputation. But I have not dwelt on the example of VNRs not just to show that the news channels airing them did not *deserve* to be accorded a reputation for “honesty” and “objectivity.” The implication I want to make is stronger than that. Just as I have argued was likely to be the case with fake news, easy acceptance of legacy news media's reputation as a credible and trustworthy source of information was the *very condition required for VNRs to be an effective means for duping the public*. Companies disguised advertisements for their products and causes as legacy news reports *precisely because* they anticipated that people who subscribed to legacy news organizations' reputation as “honest” and “objective” sources of information would be more likely to passively absorb their advertisements as *truth delivered to them by trustworthy and impartial authorities*, an attitude that sharply contrasts with the disposition of skepticism citizens tend to adopt when they realize they are being sold something. Such abuse is, I argue, always bound to arise when the public manifests routine reliance on testimonial sources *simply because* they enjoy an *established reputation* of credibility and trustworthiness. The effectiveness of VNRs and modes of deception like them is *conditional on* the public's willingness to simply accept a testifier's testimony on the grounds of their having established a reputation of reliability, which suggests that, conversely, its effectiveness is likely to be mitigated to the extent that the public is willing to subject such testifiers' claims to scrutiny rather than passively absorbing their claims.

The discussion in this sub-section illustrates the volatility entailed in heavy reliance on an individual or institution's enjoyment of a reputation of reliability as an indicator of their *actual* reliability. But by now it should be apparent that this is not the only important point I wish to make here. The more important point is that the *ongoing usefulness of reputation as an indicator*

375Price (2006)

of reliability is itself contingent on the willingness and ability of lay citizens to subject the reputations of testifiers itself to scrutiny in the light of their own judgment. If we acknowledge that some testifiers with a reputation for reliability do not deserve that reputation, that some who currently deserve it will not in the future (and vice-versa), and that even testifiers who deserve a general reputation of reliability sometimes err, there would appear to be a *standing need* for the public to *consistently* evaluate and re-evaluate the reputation of *all* testifiers. How else could we hope for citizens to decide whether the current instance is one where a normally reliable testifier ought not be relied upon? How else is the public to improve, over time, its notions of which testifiers deserve a reputation of reliability, whether that be discovering previously trustworthy ones who have taken a turn for the worse, previously untrustworthy ones whose performance has improved, or assessing the reliability of testifiers with whose reputations we have no prior familiarity? And if the public were *not* do these things, making sure our working notions of reputation are accurate and improving them over time, why should we suppose that reputation could ever come to be a useful epistemic tool *at all*? Not only does very idea *that* reputation *is* a reliable tool for evaluating the reliability of testimonial claims presuppose a certain modest optimism about the judgment of the public, but the *ongoing usefulness* of reputation as a means for evaluating testimony and testifiers *depends upon* lay citizens' willingness to subject the reputations of the testifiers it encounters *to* the scrutiny of their own judgment. But far from being actively *encouraged*, the public's willingness to subject the reputations of testimonial sources to ongoing scrutiny in the light of their own judgment is *discouraged*, even *derided*, by Progressives. For instead of accusing the citizens who have turned away from legacy news sources of having *judged incorrectly* and attempting to show them the error of their ways, Progressives treat the *very phenomenon* of turning away from sources who enjoy an established reputation of reliability as a sign of epistemic vice, society-wide “information disorder,”³⁷⁶ “post-truth malaise,”³⁷⁷ and a series of “aggressive attacks on basic Enlightenment premises.”³⁷⁸

This last comment, provided by Peter Dahlgren, illustrates the curious mingling of ideas in Progressives' opposition to the post-truth movement. Dahlgren is far from alone in condemning the post-truth movement for abandoning the ideals of the Enlightenment and

³⁷⁶Wardle and Derakhshan (2017)

³⁷⁷Lewandowsky et al (2017)

³⁷⁸Dahlgren (2018), 25

implying that their own position is more in line with them.³⁷⁹ I disagree on both fronts. Two key signatures of Enlightenment thinking were (1) a modest optimism about human judgment and (2) the idea that people ought to think for themselves rather than simply defer to established authorities – what I have called *doxastic self-rule*. Immanuel Kant's essay “An Answer to the Question: 'What is Enlightenment'” illustrates both these attributes perfectly. For Kant, the “motto of the enlightenment” was “*Sapere aude!* Have courage to use your *own* understanding!” and all that was necessary to ensure the public's eventual enlightenment was for all citizens to be given “freedom to make *public use* of [their] reason in all matters.”³⁸⁰ But it is precisely *for* having the courage to use their own understanding, and on the strength of that courage refusing to defer to “media” and “other institutions of knowledge production and expertise: universities, scientific research, schools, and the courts” that Dahlgren condemns post-truthers. At the same time, he characterizes deference *to* the institutions he names as constitutive of using “factual evidence and reasoned analysis” to draw conclusions about the world.³⁸¹ By condemning citizens for relying on their own judgment to form beliefs about the world and equating deference to established institutions as the equivalent of “reasoned analysis,” it is *Dahlgren* who violates basic Enlightenment premises. The behaviors he condemns in post-truthers are not *attacks upon* but 21st century *manifestations of* Enlightenment values. Those Enlightenment values are foundational to liberalism, so it should not surprise us that Progressives who misunderstand them so profoundly also often oppose liberal public policies and models of citizenship while supposing themselves to be proponents of liberalism.

Dahlgren also departs from both Enlightenment and liberal values by never making any attempt to explain to the reader *why* citizens ought to routinely defer to the claims of his favored institutions of knowledge production; instead, he simply asserts *that* they should do so, and characterizes anything else as by definition a sign of unreason and epistemic vice. This is so far removed from the Enlightenment's optimism about human judgment and its encouragement of *doxastic self-rule* as to be the very opposite – a categorical endorsement of established authority and a repudiation of the judgment of the lay public. This habit of *categorically endorsing*,

379This accusation is common among Progressives. See, e.g., Lewandowsky et al (2017) and Deb, Donahue, and Glaisyer (2017).

380Kant, “What Is Enlightenment?” (2004 [1784]), 54-55, *their emphasis*

381Ibid.

without argument, certain institutions that enjoy a *reputation* for trustworthiness and rebuking and accusing of epistemic vice those members of the public who, on the strength of their own judgment, refuse to place their trust in those institutions' claims, is neither true to the ideals of the Enlightenment nor a good model for the epistemic behaviors we ought to encourage in the lay citizenry. With regard to the current discussion of reputation, it implies an attitude of uncritical acceptance of authorities who happen to currently enjoy a reputation of reliability that is unlikely to aid us in – indeed, *actively discourages* us from – discovering which testifiers deserve their reputations and which testimonies even by those who do deserve that reputation may not be entirely reliable. Thus, it undermines the very utility *of* reliance on reputation as a means of vetting testimony.

The arguments and examples in this section are meant to demonstrate that while reputation *can* be of some use in determining which testimonies and testifiers to trust, its use is always *precarious*, since it is always difficult to tell whether an individual or institution's reputation is *deserved*. This is especially the case under conditions of *weak* consilience such as those that characterize the overwhelming majority of anyone's encounters with the news. Since consultation of reputation is *precarious*, citizens should not put too much stock in the conclusions of *anyone* who relies heavily upon it to derive their conclusions about which testifiers to trust, including both themselves and “fact-checkers.” At the same time, since consultation of reputation is *useful*, citizens should not be thrown into a state of epistemic despair by the prospect of figuring out the truth under conditions of weak consilience. They should accordingly *both* recognize their own limitations *and* be confident that they are as equipped as anyone else to use the tool of reputation to navigate testimony in the complicated world of the 21st century.

Section Three: Worldview

A second tool we may use to decide whether and how much to trust testimony in the absence of strong testimonial consilience is to see whether that testimony is *compatible with our view of the world*, broadly conceived. Although most of the news takes place far away from us and often in places very different from where we live, the world that contains both us and those

events is still the *same* world. It is governed by the same laws of nature and populated by human beings which, though the differences between them are so numerous and profound that I have expressed skepticism about anyone acquiring sufficient understanding of them to be counted an “expert,” nevertheless seem to admit of certain universal, or near-universal, commonalities. Our experience in the world can allow us to form some idea of the laws of nature and the tendencies of humankind, and we can use this to decide when to take some possibilities seriously and not put too much stock in others. A more conservative version of this approach would be to only invoke our worldview as a grounds for belief (or disbelief, depending on the scenario) when we find ourselves confronted by a bit of testimony we cannot, given the beliefs about which we are most confident, imagine could possibly be true (or false). In that case we can follow the advice of David Hume and cite our past “firm and unalterable experience” to reject the possibility that the testimony at hand could be true (or false).

Hume himself advocated using one's worldview to vet testimony, arguing that since “experience” is “our only guide in reasoning concerning matters of fact,”³⁸² we ought to condition our belief in testimony on our view of the world, accepting extraordinary claims only in the event that “the testimony be of such a kind that its falsehood would be more miraculous than the fact which it endeavors to establish.”³⁸³ The view of empiricists like Hume, who saw knowledge as fundamentally and exclusively gained by experience, each experience interacting with, updating and being vetted by one's prior experience, was developed and improved in the past century by pragmatists such as John Dewey, William James and, more recently, Richard Rorty. Though reliance on one's antecedent worldview to vet candidate beliefs has sometimes been criticized as embracing a myopic and solipsistic approach to belief formation (or, more cynically, belief *preservation*), I believe Rorty defended against these charges well when he argued that “people can rationally change their beliefs and desires only by holding most of those beliefs and desires constant.”³⁸⁴ Elizabeth Fricker makes a similar point in defending worldview as a legitimate tool for helping a reductionist vet expert testimony: “Fortunately we all have some basic cognitive equipment to help us assess both the sincerity and competence of others in many, though by no means all circumstances. This is because we are all experts (though of

382David Hume, *An Inquiry Concerning Human Understanding* (Bobbs-Merrill: Indianapolis, 1955 [1748]), pp. 118

383Ibid., 123

384Rorty (1991), 212

varying degrees of skill) in one special topic, namely that of folk psychology.”³⁸⁵ Both Rorty's and Fricker's arguments boil down to an acknowledgment of the inescapable role played by our extant worldview, part of which is our conception of what Fricker called “folk psychology,” in order to evaluate the likely truth or falsehood of propositions, including but not limited to those brought to us via testimony. It seems to me the pragmatists, along with Hume and Fricker, got at least this part of epistemology right: we must vet any new experience against a worldview that has been formed through many other experiences, and if the two conflict we must adapt one, the other, or (probably preferably) both. This means vetting *testimony* against our *worldview* and our *worldview* against incoming *testimony*, and using sober and considered judgment to condition each in light of the other's evidentiary clarity.

All acts of judgment entail consultation of one's worldview, so we *must* use our worldview if we are to judge at all. Still, necessary and laudatory as this belief-evaluation procedure may be, just as was the case with reputation, use of one's extant worldview to decide which testimonies and testifiers to trust comes with clear dangers. Even were we to be extremely conservative and refuse to make judgments on such matters except in those cases when one possibility seems impossible, that would still require us to first undertake the formidable task of defining what ought to count as a “impossible.” Now, many will suppose that an “impossibility” is something that is beyond the laws of nature, but the assumption of chronic fallibility shows that this cannot be the criterion we mere mortals use to assess whether something ought to count as a “impossible.”³⁸⁶ When fallible creatures like us call something “impossible,” it just means we have encountered something that we *cannot make sense of by the lights of our current worldview*.³⁸⁷ The obvious problem is that sometimes, even often, our worldview is in error. Were we to too readily write off propositions on the grounds of their being “impossible,” we would be closing ourselves off to any possibility of expanding and correcting our own worldview.

385Fricker (2006), 242

386The “laws of physics,” insofar as we know them, are not the inviolable and eternal things we generally suppose them to be. In the pragmatist vein, I suspect that the regularity with which these “laws” are tweaked, updated, and at times overthrown reveals them to be more *working models* than instances of revealed, “objective” truth. Even in physics, then, it may not be wise to respond to things that appear to be “miracles” by rejecting them outright.

387This is in fact Hume's explicit definition of “miracles” in his famous chapter on the subject (Hume 1955 [1748], 122-3).

There is an additional consideration that makes worldview an even more precarious tool for the vetting of testimony, perhaps especially in the arena of politics. This is the consideration of cognitive biases. It is clear from recent insights in psychology, and should be clear from any attentive individual's personal "expertise" in *folk* psychology, that people's thinking is often plagued by cognitive distortions, one of whose most consistent functions seems to be none other than to serve to *protect and insulate their extant worldview by distracting them from and/or skewing their perceptions of those arguments and events that imply that extant worldview to be in error*. Phenomena such as selective perception, confirmation bias, and other species of motivated reasoning are well established such biases. Unsurprisingly, these biases are frequently cited by Progressives pessimistic about the judgment of lay citizens, many of whom argue those biases to interact viciously with the media environment of the Information Age "in which ideologues create and partisan media relay compelling but misleading constructions [of fact] to the like-minded."³⁸⁸ Not only, then, is it possible that people who feel too free to use their worldview to conclude which testimonies and testifiers are reliable may erroneously discard new information that conflicts with their worldview, but given the prevalence of cognitively biased reasoning, this unhappy outcome seems rather *more* than *less* likely. And indeed this seems to be exactly what Progressives believe to be happening in the current informational environment, which they claim is filled with political "filter bubbles"³⁸⁹ and "echo chambers," in which cognitive biases like those mentioned above run rampant and contribute to dangerously growing levels of polarization.³⁹⁰

I will have more to say about these issues in the next section on perspectival diversity, but for now I only wish to emphasize two things. The first thing I want to emphasize is that I believe the role played by cognitive biases in distorting human judgment to obviously be significant. Cognitive biases are *very* likely to distort *anyone's* filtration of new information in such a way that it excessively and undeservedly confirms their worldview, and to some degree we of the contemporary era have witnessed this tendency's unfortunate synergy with the onset of Information Age and its sudden provision of the ability to choose from many different available

³⁸⁸Jamieson (2015), 69

³⁸⁹Pariser (2011)

³⁹⁰Sunstein (2017). In agreement with numerous empirical scholars of political communication (e.g. Garrett 2017, Flaxman et al 2016), I believe these worries about "filter bubbles" and "echo chambers" to be dramatically overblown, but will save that argument for later in the dissertation.

data sets and news reports to corroborate one's own pre-existing view of reality. When we assess news agencies and the reports they disseminate, *all* of us are likely to reason in a way that inclines us to trust some sources and to distrust others in a way that skews excessively in favor of our pre-existing worldview. Indeed, the generality of judgmental realms over which cognitive biases seem to regularly distort our reasoning is difficult to overstate. Such biases seem regularly to affect people's judgments even in seemingly “objective” and uncontroversial epistemic realms such as mathematics.³⁹¹ They are common among both people who have received high levels of education, including experts within their domain of expertise, and those who have not.³⁹² They are common among those who are measured as having high levels of general intelligence and those who do not.³⁹³ They are common all along the political spectrum.³⁹⁴ Indeed, although “[i]ndividual differences in cognitive processes and susceptibility to [cognitive biases] are not uniformly distributed,” meaning that not *everybody* manifests the *same* propensity to fall prey to such biases, those differences are “understood to be the product of *individual* or *selective* contexts.”³⁹⁵ In other words, although we have good reason to suppose that *some* people are better at combating their cognitive biases, we cannot tell by any easily identifiable sign, such as social position, educational attainment, or other obvious sign, who these people are going to be.

While consideration of cognitive biases gives some support to Progressives' pessimistic views about the judgment of the lay public, it undermines the Progressive notion that any identifiable individual or group, be it the legacy news media, a government agency, or “fact-checkers” employed by Facebook or Google, ought to be treated as if their judgment is likely to be consistently and reliably superior to the judgment of lay citizens. Our inability to reliably identify those individuals or groups who are especially likely to resist their cognitive biases and judge in a more neutral or “objective” manner reinforces the argument presented in the previous section citing the *absence of identifiable expertise about the news* as a reason why citizens should not feel compelled to defer to any particular individual or institution's views of what news they ought to place their trust in, and certainly ought not to accept the imposition of any sort of authority to *cancel* what news they can or cannot see. Citizens have every reason to suspect that,

391Nurse and Grant (2020)

392Tversky and Kahneman (1974)

393Teovanovic et al (2017)

394E.g. Taber and Lodge (2006), Nyhan and Reiffler (2010)

395Collins and Laughran (2017, 6)

even if we made the (reckless) assumption that such authorities would act with the public's best interest at heart, they would be subject to the same sorts of biases in judgment as those to which lay citizens themselves are subject. This gives lay citizens additional reason to feel uncompelled to defer to any individual or group's claims about which news reports and sources ought to be trusted rather than relying on their own judgment, and to resist any effort to censor its information on the grounds that the citizenry is less fit to judge which testimonies and testifiers are reliable than whatever group is doing the censoring.

The second point I want to emphasize about the use of *worldview* to evaluate testimony runs in a more optimistic direction. This point is that, even as we acknowledge the *existence* and *significance* of the worldview-insulating problem of cognitive biases, and accordingly ought to be deeply suspicious of any claims of expertise that imply immunity from such biases, we should not allow that acknowledgment to lead us to the level of deep pessimism about human judgment that is unfortunately in vogue among a certain vein of social scientists and public commentators at the moment. Many of the accounts written by these figures at present seem to imply that democratic citizens are not now, were not before, and are not capable of ever being *rational* in any meaningful respect.³⁹⁶ Yet, while cognitive biases *do* naturally distort our reasoning and *can* cause serious problems if left to run amok, and while as yet we cannot reliably identify any particular group of people who suffer from them appreciably less than the average citizen, we have good reason to believe their influence can be *mitigated* by encouraging *more critical and active evaluative habits* on the part of individuals.³⁹⁷ As William Minozzi and Michael Neblo put it, “the evidence for affectively 'tainted' reasoning hardly constitutes grounds for despair about the role of reason in democratic politics.”³⁹⁸ Instead of despairing about the role of reason in democratic politics, our take-away from the cognitive biases literature ought to be that we should be *encouraging* the lay citizenry to *more actively and critically evaluate* the testimonial claims

³⁹⁶Excessive cynicism about human rationality is so common in the social sciences it seems hardly worth citing specifics. It is characteristic of a large chunk of the “heuristics and biases” tradition inspired by the work of Daniel Kahneman and Amos Tversky, which is excellent empirically but, I believe, rather lacking when it comes to interpreting its meaning for individuals, society and politics. For two canonical examples of what I take to be excessive cynicism derived in part from the motivated reasoning literature in social science, see Lodge and Taber (2013) and Achen and Bartels (2016).

³⁹⁷For an excellent recent defense of a similar position, with numerous references to support in the empirical literature, see Arceneaux and Vander Wielen (2017)

³⁹⁸Minozzi and Neblo (2019), 15

they encounter in the informational environment, rather than passively cruising through that environment relying on the “fast” cognitions and intuition-friendly heuristics³⁹⁹ that most commonly lead to bias.

It is true, of course, that such active and critical evaluation of testimony entails consulting one's antecedent worldview because *all* judgments necessarily do – as Rorty put it in the passage cited above, we can only rationally change our beliefs by holding most of our other beliefs constant. And of course, since our worldview may be in error, using it as the standard of evaluation may lead us astray. But that danger is endemic to *all* human thought. The solution cannot be to *abandon* the very *use* of our worldview as a means for evaluating candidate beliefs, for that would entail the very abandonment of judgment, which, flawed though it may be, is our only means for actually *evaluating* candidate beliefs rather than adopting them on an epistemically arbitrary basis. Nor, since we have no good reason to suppose the worldview of any identifiable individual or group to be more accurate than the lay citizen's or their judgments to be less subject to bias across all the epistemic domains encompassed by the news, should we advocate for citizens to offload their burden of judgment onto any group of would-be authorities. Instead, the proper approach should be to *encourage* lay citizens to *more actively and consistently* exercise their judgment, along with – as I will argue in the final chapter – educating them *about* the dangers of cognitive biases and try to motivate them to combat them. Again, however, this *critical and active* exercise of judgment on the part of lay citizens is precisely what is *discouraged and disparaged* by Progressives. It is discouraged by Progressives' suggestions that citizens ought to routinely defer to the testimonial claims of a pre-defined set of institutions, rather than predicating their trust on a critical and active exercise of their own judgment. And it is routinely disparaged by Progressives whenever post-truthers *do* cite their own judgment, rooted in their extant worldview, as a grounds for refusing to defer to Progressives' favored authorities' claims. As we have seen, Progressives regularly condemn such behavior as definitionally irrational, solipsistic and dangerous to democracy.

In sum, just as with reputation, consideration of use of one's worldview as a means of vetting testimony under conditions of weak consilience leads to the recommendation of what may at first glance seem a curious combination of epistemic modesty and assertiveness on the

399Kahneman (2011)

part of the lay citizenry. Lay citizens' epistemic *modesty* should come from an acknowledgment that, while consultation with one's worldview *can* be useful in helping us adopt more reliable testimonial claims and avoid placing our belief in those that are less reliable, its use is *precarious* because everyone's worldview is imperfect and our thinking is often skewed by cognitive biases that distort our reasoning. Thus, everyone's confidence in the beliefs they form in their beliefs about which news reports and sources to trust based on that worldview ought to be *limited*. Yet, these very same considerations simultaneously imply support for a certain form of epistemic *assertiveness* on the part of citizens. The same limitations that inevitably limit the confidence any individual citizen ought to place in the beliefs they adopt about the reliability of testimonies and testifiers ought also to apply *generally*, for we have no reliable means of deciphering which groups' or individuals' worldview is more “correct” or whose cognitions are less skewed by bias. Citizens ought not conclude from their own imperfections in judgment in the arena of testimony that they ought therefore to *defer* to the claims of any set of would-be authorities, for no argument can be made identifying any particular such authority as an “expert” in the use of their worldview as a means of evaluating testimony under conditions of weak consilience in a way that ought to be accepted by lay citizens. Therefore, instead of conceptualizing lay citizens as incurably riddled by cognitive biases and fostering the attitude of routine deference to authorities like those in the legacy news media in the way commonly done by Progressives, we ought to encourage citizens to be *more active and critical* in their evaluations of testimony and testifiers and educate them about the dangers of cognitive biases and the tools and techniques available to *all* of us – not just some supposed group of cognitive “experts” – to combat them.

Section Four: Perspectival Diversity

A final tool available to the knowledge-seeker trying to get the best out of testimony, which I believe to be not only valuable in itself but also crucial in getting the best out of the other tools I have mentioned, is seeking out *perspectival diversity*. This means simply that one ought to seek as many *independent* sources of testimony as one can before incorporating a claim into one's stock of beliefs, where *independent* just means *not systematically connected in any way that seems indicative of or likely to impose a uniform point of view*. The value of this sort of

perspectival diversity stems from exactly the same considerations as those I explained in defending the value of consilience. That is, *all* of our knowledge, and all our *means* of acquiring knowledge, are subject to error, which makes hinging our beliefs on a single grounds precarious. While this does mean that any *particular* belief we adopt, and any *grounds for belief* we employ, may be in error, it is less likely that multiple *independent* grounds will all be flawed in a way that skews the beliefs we derive from them in an identical direction. Therefore, we should be more confident in beliefs that have multiple, independent grounds than in those that have only a single grounds. I presented this argument at the beginning of this chapter to defend consilience broadly conceived, or what I called *strong* consilience. But even when we are under conditions of *weak* consilience, the same considerations imply that more diversity is better than less. That is, the more independent sources of testimony we encounter that all endorse a given belief, the more confidence, *ceteris paribus*, that belief warrants.

The definition provided above implies that there are several ways *independence* between sources may commonly be violated in a way that ought to limit our readiness to place extra confidence in certain testimonies, even when those testimonies disseminate from multiple sources. Independence can be violated *either/both* by (a) the multiple testimonial sources from which a claim disseminates *already possessing a similar point of view* and/or by (b) testimonies going through a *homogenizing process of articulation and dissemination*. The possession of a similar point of view from the beginning violates independence because people with a similar point of view are, *ceteris paribus*, likely to notice similar things about a situation, interpret those things in a similar light, and describe it in similar terms. These similarities mean that whatever flaws are endemic to that point of view are comparatively likely to skew *each* of the testimonies disseminated by *each* of the testifiers *in the same way*, which necessarily ought to lessen our confidence that the testimonial account that eventually reaches the public is an accurate and complete depiction of the *full story*. Likewise, even testimonial accounts that *originally* disseminated from *multiple* sources with *different* points of view can lose their independence by going through a *homogenizing process of articulation and dissemination*, such, for example, as requiring the testimony to be expressed in a given linguistic style, emphasizing certain features of the story that fit an established and familiar range of themes, and selecting and omitting aspects of the story based on a given set of tastes, sense of propriety, and other sorts of values.

Since the value of consilience, and therefore perspectival diversity, is proportional to the diversity of our grounds for belief, which makes those beliefs less precarious by making us less dependent on any single grounds for them, any process that *reduces* that diversity necessarily also reduces consilience, and therefore ought, *ceteris paribus*, to temper our confidence in beliefs that rely solely or predominantly on testimonies that have gone through such a homogenizing process of articulation and dissemination.

The institutions favored by Progressives, such as the “scientific community” and the legacy news media, sometimes enact procedures that may at first seem to acknowledge the value of and implement perspectival diversity into their knowledge production and dissemination processes. For instance, the academic publication process entails anonymous peer-review, part of whose purpose is (in theory) to run the ideas of a single writer or group of writers, who are coming at things from one perspective, by reviewers anticipated to be coming at things from a perspective that is at least slightly different – for instance, because they subscribe to a slightly different theoretical paradigm in their field, or because their research incentives are different from those who wrote the study they are reviewing. Similarly, as we shall see, gathering information from multiple sources and subjecting claims to a fact-checking process conducted by someone other than the reporter writing the story are supposed to be fundamentals of the process of professional newsmaking. Requiring *multiple* sources would seem to increase the odds of bringing some perspectival diversity to bear on a story, and fact-checking by someone other than the story writer herself would seem also to be a potentially valuable means of injecting some perspectival diversity into the publication process. Indeed, all these processes, both by members of the “scientific community” and by professional journalists, are mentioned by epistemologist Alvin Goldman in articulating a philosophical defense of the Progressive position endorsing the superiority of the legacy news media to online-era alternatives and the importance to democracy of the public's willingness to trust the former more readily than the latter. Yet, tellingly, Goldman himself does *not* attribute the value of these procedures to notions of consilience, perspectival diversity, or anything related to them. Quite the opposite: Goldman's account implies that a *single* perspective and a *single* set of *already-established* methods – the perspective of the professional journalist and the methods employed by the legacy news media – already exists, and that we should be confident that this perspective and these methods are bound to produce and

disseminate more reliable knowledge than whatever might make its appearance on the blogosphere. Far from emphasizing the value of *perspectival diversity* due to *chronic fallibilism*, Goldman's account emphasizes the value of *professionalism* and *expertise* and appears to suppose that we need not worry too much about fallibilism when it comes to those professionals and experts. By looking at Goldman's essay in detail, as well as the arguments of Richard Posner to which Goldman's own article is framed as a response, we can see the ways that, even while it might seem to endorse institutions whose design implements measures to increase perspectival diversity, the position of Goldman and Progressives actually tends to *undervalue* and often *actively oppose* perspectival diversity, largely because it underestimates the scope of chronic fallibilism, imputing it only to lay citizens and not the established authorities favored by Progressives. We shall see that the account of Goldman betrays a deep pessimism about lay citizens' judgment and encourages those citizens to adopt an attitude of excessive credulity toward and routine deference to certain established authorities in the same way I have argued to be characteristic of the Progressive approach to civic epistemology.

Goldman's essay, titled "The Social Epistemology of Blogging," was written as a rebuttal to an article written by judge and legal theorist Richard Posner, who expressed a comparatively friendly and optimistic attitude toward the "blogosphere" vis-a-vis the legacy news media. Posner's original argument in favor of the blogosphere implies the value of perspectival diversity in multiple respects. For instance, Posner approvingly notes the sheer number and variety of bloggers, many of whom "specialize" in different subject areas, allowing "unorthodox views to get a hearing." He also emphasizes the "better error-correction machinery" of open-forum comments allowing anyone to publicly opine about the veracity of the bloggers' testimonies in a way that contrasts starkly the mainstream news, which has no such feedback mechanism and is itself in charge of deciding whether or not a retraction – of its *own* claims – ever gets printed at all. Posner acknowledges that professionalism and a curated publication process have some potential benefits but argues that, once "all the pluses and minuses of the impact of technological and economic change" on the informational environment "are toted up and compared, maybe there isn't much to fret about."⁴⁰⁰ The gist of Posner's argument seems to be to put the blogosphere and legacy news on roughly equal epistemic footing – *equal* but not *equivalent*. The

⁴⁰⁰Posner (2005, 10-12)

blogosphere, for Posner, is simply a *different* means of generating and disseminating information than the legacy news media, one that comes with its own unique set of advantages and drawbacks. Its main advantages are the diversity of perspectives it can bring to bear on political questions and its dynamic and open-source means of facilitating contestation of posters' claims by outside sources through "comments" boxes and their like. Meanwhile, its main disadvantage is that, since the blogosphere does not have the same ethos of balance and objectivity as professional journalism, bloggers can be brazenly partial in their reports and motivated readers can fuel their biases by selecting their information on the basis of their ideology rather than truth. While acknowledging that the blogosphere comes with the sorts of disadvantages attributed to the 21st century informational environment emphasized by Progressives, then, Posner's argument is that these disadvantages may well be made up for by the blogosphere's considerable superiority when it comes to facilitating *perspectival diversity*.

Goldman disagrees with Posner's analysis, defending the Progressive view that the legacy news media is more reliable epistemically than the blogosphere and that citizens trusting in it is of special importance to democracy. Before exploring Goldman's reasoning in detail, it is worth starting off by noting that Posner's original defense of the unfiltered informational environment of the internet did *not* lead him to the conclusion that the blogosphere was *superior* to the legacy news media or that the unfiltered informational environment provided by the internet was *ideal*; instead, Posner's conclusions are quite modest. "[P]robably there is little harm and some good in unfiltered media," he writes in one passage before later reaching the aforementioned conclusion that "maybe there isn't much to fret about."⁴⁰¹ Goldman himself acknowledges the modesty of Posner's claims, characterizing Posner's argument as amounting to the claim that "blogging, as a medium of political communication and deliberation, is no worse from the standpoint of public knowledge than conventional journalism." It is clear that Goldman thinks we have grounds to reject even such a comparatively modest position. Goldman's argument clearly implies that the legacy news media is *largely reliable* as a source of information and that it is *more reliable than* the blogosphere, as he describes himself as providing "an argument for the superiority of conventional news media over blogging, so long as knowledge and error avoidance are the ends being considered." The key notion around which Goldman frames his defense of the legacy news

⁴⁰¹Ibid., 12

media vis-a-vis the blogosphere is the notion of *filtering*, which he defines as occurring when “an individual or group with the power to select which...messages are sent via [a] designated channel and which are not. When a gatekeeper disallows a proffered message, this is filtering.” In other words, *filtering* means a process by which some set of authorities determines which testimonies, and aspects thereof, are *allowed* to be disseminated by a given institution. Goldman's argument seems to be that *as a general rule* information that goes through processes of filtration ought to be considered more reliable than that which does not, or, as Goldman puts it, “from a veritistic point of view, filtering looks promising indeed.”⁴⁰² To illustrate the epistemic benefits of filtering, Goldman uses the examples of peer-review practiced by conventional scientific journals and the rules of evidence and procedure that govern criminal trials, arguing that the filtration processes employed in these scenarios is precisely what increases the epistemic reliability of the testimonies and decisions that go through those processes – in this case, academic articles and trial verdicts. He then mentions several processes the legacy news media uses to filter the information it publishes – specifically, filtering out “extreme views,” employing fact-checkers “to vet a reporter's article before it is published,” requiring more than a single source before publishing a story, and limiting use of anonymous sources – before finally concluding that, “from a veritistic point of view, filtering looks promising indeed.” He takes this line of argument to demonstrate “the superiority of the conventional news media over blogging” in terms of “knowledge and error avoidance.”⁴⁰³

Goldman's essay is valuable in that it is one of the exceedingly few examples I can find of anyone attempting to actually supply a rigorous defense of Progressives' veneration of the legacy news media vis-a-vis online-era informational sources, as well as of the oft-repeated view that democracy's well-being depends in some appreciable way on lay citizens being willing to place their trust in it. There are a number of things to point out about Goldman's response that are both typical of Progressives' response to 21st century communications technologies and, in my view, problematic. The first thing to point out is Goldman's veneration of, and credulity toward, established authority, and his suspicion of and antipathy toward non-mainstream alternatives. In the passages cited above, and throughout the essay, Goldman attaches the adjectives

402Goldman (2010), 116-117

403Ibid., 113-115

“conventional” and “traditional” to the knowledge-disseminating authorities he favors while routinely disparaging the newer sources made available by the internet. He also takes it for granted, as do many Progressives, that one of the primary dangers of the internet is that “[b]ad people” tend to “find one another in cyberspace and gain confidence in their crazy ideas,” and that therefore a “virtue of conventional media...that they filter out extreme views.”⁴⁰⁴ As I have already shown on numerous occasions, this reflexive affinity toward established authority combined with suspicion of and antipathy toward digital-era newcomers is typical of the Progressive position. So, too, is the casual lumping-together of the concepts “crazy,” “bad” and “extreme,” and the implicit acceptance of the idea that “extremists” are easily identifiable, categorically unreasonable, and dangerous to democracy.⁴⁰⁵

I believe all these tendencies manifest in Goldman's essay and typical of the Progressive approach to the 21st century informational environment are detrimental to liberal democracy, both in terms of improving our *collective knowledge*, which appears to be Goldman's main concern, as well as in terms of improving *political outcomes*, which Goldman also cites as a motivating force for his work. The approach of Goldman and Progressives encourages a reflexive resistance to unfamiliar (“extreme”) ideas and lifestyles that is likely to both deter the growth of knowledge and to be used to disparage and repress individuals and groups who subscribe to ideas and lifestyles that cannot be described with Goldman's favored adjectives of “traditional” or “conventional.” Like many other liberals, I take the propensity to resist challenges to convention, both in terms of established authority and in terms of the range of views a society views to be “acceptable,” to have been one of the most pervasive and dominant tendencies of societies around the world, and to have been one of the most consistent obstacles to growth both in terms of human knowledge and in terms of political outcomes through the ages. Since societies around the globe in all ages have demonstrated a consistent and apparently natural tendency toward excess in their preference for preserving convention and repudiating challenges to it, they need no further encouragement in that regard. In fact, I take one of the most valuable aspects of the Enlightenment-inspired liberalism Progressives sometimes claim themselves to be the standard-bearers of as having recognized that one of the chief *vices* in

⁴⁰⁴Ibid.

⁴⁰⁵This tendency is especially rampant in the various writings of Cass Sunstein, who is one of the most prolific defenders of the Progressive suite of positions (e.g. Sunstein 2009, 2014, 2017).

human history is people's readiness to embrace, without much question, whatever authorities and conventions they inherit as the "legitimate" ones and to characterize challengers of those authorities and flouters of those conventions as lunatics, heretics and infidels. Since the natural tendency of humanity is to *too easily* embrace these authorities and *too readily* ignore, ridicule and abuse challengers to them, we ought as a general rule to try to intentionally cultivate the *opposite* habits in citizens, so as to counterbalance this natural tendency. Indeed, I believe this may be the single most precious legacy handed down to us by the Enlightenment and by liberalism as exemplified in the work of Kant and Mill. But this is the *opposite* of the approach taken by Progressives, who typically demand deference to "established," "conventional" and "traditional" authorities, often (Goldman accepted) without even attempting to provide a reasoned argument in favor of that demand, and presume any decision not to is an example of epistemic vice. It is likewise the opposite of the attitude implied by the easy lumping-together of the terms "crazy," "extreme" and "bad" so often manifest in the writings of Progressives, and in their insistence that the ability for "extremist groups" to form like-minded communities online is bad for democracy. That view implies that we as a society have already figured out which points of view and lifestyle are *truly* acceptable, and that our finding certain positions "extreme" is guaranteed to be an indicator of that group's evil or misguided nature, rather than an indicator of widespread societal prejudice and excessive comfort in our own established views.

Yet, history, both distant and recent, gives us much reason to doubt whether our current notions of acceptability are as enlightened as this view implies. Consider just a few examples from the last decade: the Black Lives Matter movement, the conversation about alternative gender identities, and our evolving views about the morality and legality of drug use. All these conversations involve groups of people and the exchange of ideas that many in the recent past, and indeed many in this very moment, would consider "extreme" and "dangerous to society." Yet, I argue that in all these cases, as in many others, these "extremist" groups are responsible for improving our collective beliefs and expanding our once-too-narrow conceptions about what sorts of behavior and beliefs ought to be considered "acceptable." I also strongly suspect that 21st century technologies have played a pivotal role in strengthening these groups and allowing them to get their views aired. The Black Lives Matter movement arose in the social media era and has always relied largely on social media as an organizing tool and on videos captured by mobile

phones and posted on public platforms like Youtube to air videos to demonstrate aspects of the reality of police relations with Blacks to an audience composed of many people who do not themselves experience such relations. Likewise, alternative gender communities have long used Tumblr and other internet-era tools to find and “follow” one another, hear each other's stories, gain a feeling of belonging and acceptance, and develop their views. Although gender non-conformists have long existed, it is only now, in the digital era, that their existence is beginning to be noted by the public and the issues that face them addressed, both legally and culturally. Likewise, users of marijuana and other recreational drugs use many online forums to share information about the safety and appropriate doses of various drugs, organize support for the legalization of certain taboo substances, and quickly spread news of legal progress in certain jurisdictions that legalize or de-criminalize drug use, as has increasingly been the trend in American cities and was recently done at the state level in Oregon.

In all these cases, as well as numerous others that could be mentioned, I argue the ability of groups conceived of as “extremists” by appreciable segments of society to find each other online, motivate one another by lending each other mutual support, and gain public awareness of their movements, has improved American democracy both epistemically and politically in the internet age whereas very little and slow progress was being made in those same issue areas in the pre-internet era. While the digital age allowed these groups to form, sustain, organize, and get their views aired, it is very likely that the pre-internet age not only did not facilitate but often actively deterred such efforts. The legacy news media that dominated the pre-internet era was largely comprised of white journalists from middle class backgrounds, trained in the same sorts of journalistic traditions and working for institutions whose economic motivation demanded them to market their news toward the “average” American. Being drawn so regularly from similar demographics and economic stations, legacy journalists likely already shared similar – and similarly flawed – understandings of what views and communities ought to count as “extreme.” Undergoing largely uniform training at institutions designed and dominated by people coming from similar backgrounds, this original tendency would be likely to be reinforced. And being employed by large media companies attempting to appeal to a broad, mainstream audience, financial incentives likely pushed legacy news to avoid airing the views of “extremist” groups at all, and when they did air them, likely deterred them from airing them in a sympathetic

light.⁴⁰⁶ In other words, the composition and publication process employed by legacy news companies *deterred* the same level of progress we now see in the internet era with regard to public awareness and prioritization of police brutality, gender, and drug issues *precisely because* it was resistant to “extreme” views and communities, often ignoring them and, when it did give the spotlight to such views and communities at all, depicting them in the same light as that in which they were viewed by the journalists themselves and society at large – that is, as “crazy,” “bad,” “extreme,” and dangerous to society. Goldman and Progressives' tendency to lump together these same adjectives and to champion an informational environment dominated by this same type of news media shows that they underrate the extent to which the legacy news of the past ignored and demonized groups with legitimate political grievances *by* labeling them “extremists.” Were citizens to adopt Goldman and Progressives' attitude of reflexive opposition to those groups and turn away from the blogosphere and social media in favor of the legacy news media in the way Goldman and Progressives advise, our epistemic and social progress when it comes to expanding our understandings of the different communities of people with whom we share a society and revealing and upending the prejudices that are as likely to prevail in our current society as those we know to have prevailed in those previous would be diminished.

At the same time as it feeds the already-too-easy tendency for peoples to oppose challenges to convention and label those who flout it as “extremists” who pose a threat to society, the position of Goldman and Progressives demonstrates and encourages an attitude of excessive credulity toward and acceptance of their favored institutions as authorities. That is, they do not subject their favored authorities to the same degree of scrutiny to which they advocate for citizens to subject newer informational sources that have emerged in the online era. Goldman's depictions of the legacy newsmaking and scientific publications processes are a case-in-point. On the legacy newsmaking front, Goldman cites the employment of fact-checkers to “vet a reporter's article before it is published,” the requirement of a story to cite “more than a single source” in order to qualify for publication, and imposing limits on reporters' reliance on anonymous sources as examples of “filtering techniques” legacy news organizations employ to

⁴⁰⁶Indeed, this is widely acknowledged to have been the case in the era of the Fairness Doctrine, which required media companies in the broadcast era to give equal time to opposite political points of view in a way very similar to the “both sides” ethos long championed by professional journalists. In actuality, the Fairness Doctrine appears to have worked as a “tax on controversial speech,” deterring legacy news channels from airing politically charged programming, especially “controversial” points of view (Hazlett and Sosa 1997, Sunstein 2017).

improve the “veritistic quality of their readers’ resultant beliefs.”⁴⁰⁷ As for “conventional scientific journals,” Goldman mentions peer review as a filtering process that, again, increases the veritistic quality of the beliefs of those who place their trust in such articles.⁴⁰⁸ While I will not deny that all of these processes *sound* like good ways of improving the reliability of the claims disseminated by their respective institutions, there is a crippling weakness to Goldman’s article that should not be overlooked, and which demonstrates why I believe the Progressive point of view is so dangerous. The weakness of Goldman’s position is that the value of the processes he mentions comes not from their having been stipulated for the purposes of argument, as Goldman does, but instead from their *actually* and *responsibly* being practiced by the authorities (supposedly) in charge of carrying them out. As epistemologist David Coady writes in a critique of Goldman’s position, “[t]his is, to a large extent, an empirical issue,” so it is “striking...that Goldman does not refer to any empirical work on media performance.”⁴⁰⁹ The same is true of Goldman’s claims about scientific peer review. Goldman simply *supposes that* the processes he mentions are being practiced regularly and responsibly by the authorities he endorses, and concludes from that supposition that such processes are beneficial epistemically and democratically, and that citizens are endangering democracy by failing to properly place their trust in them.

Is it a good idea for citizens to follow Goldman’s example and *simply suppose that* scientists are responsibly practicing peer-review and that major news organizations are employing all the techniques Goldman mentions? I do not believe it is. This is *not* because I think citizens have reason to be especially cynical about the morals of academics and professional journalists. Instead, citizens ought not to *simply suppose that* epistemic authorities are exercising their stations responsibly as a matter of basic prudence. As I argued in the last chapter, this is perfectly compatible with the attitude of modest optimism I argue to be central to defenses of liberal democracy. Even if we believe, as I do, that *most* testifiers, including journalists and scientists, are acting in good faith *most* of the time, our shared expertise in folk psychology tells us that we can count on at least *some* of them declining to do so *some* of the time, and those exceptions sometimes being of vital importance. The same sorts of folk

407Goldman (2009), 117

408Ibid., 116

409Coady (2011), 141

psychological considerations imply that the tendency to abuse or become negligent in the exercise of one's epistemic authority is likely to become increasingly common to the degree that those people we treat as epistemic authorities learn that their claims are accepted without much scrutiny, as Goldman does in his essay and as Progressives often do generally.

Despite Progressives' tendency to treat them as special cases, it is worth stressing one more time that these folk psychological concerns apply as readily in the domains of science and legacy news as in any other domain of human affairs. It is likely that most people reading this are or have been members of academia, and are therefore at least marginally familiar with the “inside” of the peer-review process. Even those with limited exposure to that world should know that, although the peer-review process is supposed to be anonymous and impartial, and although there are many truly sincere individuals who carry out their peer-reviewing duties in exemplary fashion, the high-minded ideals anonymous peer-review is meant to embody are nevertheless skirted with some degree of regularity in practice. Sometimes, this is because anonymity cannot be effectively preserved in the small fields of specialization characteristic of the upper levels of academic research, where people become familiar with each other's research, perspectives, writing styles, and so forth. Other times, it is because reviewers are biased in their treatment of prospective articles, either ideologically or because they have an interest in promoting their own career by being cited as many times as possible. It is far from unheard of for reviewers to reject articles that either do not align with their own ideological perspective or which fail to cite their own work. Nor should this surprise us, since researchers are human beings and, like other human beings, often respond to ideological, monetary and status incentives. Those incentives may, and in many cases do, lead some researchers to conduct their peer-review duties in a way that not only does not help the “scientific community” pursue better ideas, but in fact actively inhibits that process.

Similarly vulnerable to the vagaries of human behavior are the processes of filtration Goldman cites admiringly in the legacy news media are skirted with at least as much regularity. Take, for instance, fact-checking. The only attempt I am aware of that has attempted to empirically determine how often the content of news stories was *actually* fact-checked was conducted by researchers at Cardiff University in 2008. It found that “only “12% of these stories

showed evidence that the central statement had been thoroughly checked.”⁴¹⁰ Similarly dubious is Goldman's claim that the legacy news media consistently insists its stories be based on information from multiple, independent sources or limit its reliance on anonymous sources. Not only do most major news organizations rely predominantly on a single source – one of Reuters, The Press Association, or the Associated Press – for the vast majority of their stories⁴¹¹ but, as Coady convincingly illustrates, it is far from unheard of for legacy news stories of vital public importance to rely not only on a single source but indeed for that source to have a likely conflict of interest deterring them from telling the full story.⁴¹² Likewise, anyone who reads the newspaper regularly should seriously doubt whether the legacy news media puts much pressure on its reporters to limit their reliance on anonymous sources, even in cases where that anonymity seems as likely to serve to insulate a source from repercussions if they are found to be lying as it does to insulate them from being punished as a whistle-blower.⁴¹³ Again, we should not be too surprised at any of this. The behavior of legacy journalists is as likely to be influenced by the incentives held out to them as the rest of us, and again some of those incentives run against their conducting the above-mentioned duties in an ideal way. For instance, individual journalists may have an opportunity to publish a “scoop” full of claims that could not be published, or would not be as sensational, if all of them needed to be fact-checked or if certain other perspectives were brought to bear in the article. Likewise, at the institutional level, all the procedures named above cost news-making organizations time and money, and as such hold out the temptation for organizations to cut corners. It would be naive to think legacy news companies uniformly and universally resist such temptations. Again, then, while the filtration processes Goldman names, *if properly practiced*, would be excellent ways of improving the reliability of the claims of the institution whose members practiced them, we cannot simply *suppose* that such best-practices are actually being employed. On the contrary, given the incentives involved, it would be naive for citizens to simply suppose that those practices were being employed. Moreover, just as was

⁴¹⁰Davies (2008), 53

⁴¹¹*Ibid.*, 75-94

⁴¹²Coady (2011), 153-4. Davies' Cardiff study also suggests the requirement of multiple sources guides journalistic practice only sporadically at best. In that study, fully 70% of “stories which relied on a specific statement of fact...the claimed fact passed into print without any corroboration at all” (Davies 2008, 53).

⁴¹³If your own news-reading experience does not corroborate this, you may look to Coady's examples for a variety of illustrations (Coady 2011, 154-8).

the case with peer review, if *improperly* practiced, the filtration process employed by the legacy news media can actually *detract from* the reliability of the claims disseminated by an institution.

Tellingly, Goldman makes no mention of these real-world issues. Instead, he seems to accept implicitly that the members of the “scientific community” and legacy news media either possess an exemplary moral compass or are animated by incentives that push them uniformly in the direction of truth and honesty. In so doing, although he goes farther than most Progressives in at least *attempting* to articulate a defense of the position that citizens ought to trust the claims of the legacy news media more readily than those they find in the blogosphere, he still demonstrates a similarly uncritical and quiescent attitude toward these two sets of epistemic authorities as that of Progressives. It is precisely this uncritical attitude that *leads* Goldman and Progressives to so strongly prefer knowledge production and dissemination processes by *established* institutions peopled by “professionals” and “experts” and to underrate the value of the sorts of *error-correction and contestation* mechanisms emphasized by Posner as the chief advantage of the blogosphere and arguably even more characteristic of the informational environment provided by social media with its “likes” up- and down-votes, and the prominent visibility of comments.

How does all this connect to the idea of perspectival diversity? It does so in two ways. First, in arguing that the chief value of peer-review and the aforementioned processes used by the legacy news media stems from their *being processes of filtration by epistemic authorities*, Goldman makes no mention of the fact that these same methods are also means of passing claims through a process that subjects them to *multiple (theoretically) independent perspectives*. That he emphasizes the veritistic benefits of having an established authority vet claims and makes no mention whatsoever of any veritistic benefit of bringing multiple independent perspectives to bear on a situation suggests he sees little epistemic value in the latter. Second, by venerating the established authorities he does and so heavily emphasizing the benefits of filtration by them while simultaneously disparaging the non-establishment masses on the blogosphere, Goldman implies that we can *already* be sure, to a high degree of confidence, *which* authorities, institutions and procedures are the *right* ones – that is, which ones bring us the best knowledge in the most trustworthy way. Were this not the case, Goldman would surely place more value than he does in the mechanisms of error-correction and contestation highlighted by Posner, mechanisms that leave room for the discovery of both good-faith errors by established epistemic

authorities and abuses of that authority by those acting in bad faith. The position of Goldman and Progressives, then, appears to place comparatively little epistemic value in perspectival diversity, implying that we *already* know, at least in some important cases, *which* perspectives are the “right” ones and which authorities are using their power responsibly. This is why Progressives' view of responsible citizenship entails *deferring to* those authorities rather than subjecting them to critical scrutiny, and why Goldman's account emphasizes the value of *information filtering by the right authorities* and says little to nothing about the value of or need for perspectival diversity either in the internal knowledge production and dissemination processes used by his favored institutions or as a means of subjecting those institutions to critique from without.

In contrast to Goldman and Progressives, I view the seeking out of perspectival diversity to be one of the most, if not the single most, important means humans have available to them for improving the reliability of their beliefs. This applies to professionalized institutions like those of the “scientific community” and legacy news media as much as it does in all other epistemic arenas. The value of academic peer-review is not, in my view, chiefly that it entails *filtration by authority* but instead that it brings *multiple perspectives to bear* on a problem within a given epistemic domain. To be sure, those perspectives must all share expertise in a given field, but that expertise is not the *sole* or even *primary* reason peer-review is a good idea. After all, the original writers of academic articles are themselves usually experts in the same sense as those reviewing their work. The primary reason peer-review is valuable stems from the ability of people coming at something from multiple independent perspectives to spot one another's errors – both those made in good faith and those impelled by baser motives. Similarly, the primary virtue of the publication procedures employed by professional journalists is not that authority is imposed by virtue of those procedures but rather that they are means of injecting the right kinds of perspectival diversity into the journalistic story-writing process. Passing a story by fact-checkers serves the exact same role as peer-review, helping to guard against errors made in both good faith and in bad. Requiring multiple, independent sources for each story is a means of making sure that story is an accurate depiction of the situation, hedging journalists against the possibility that a single source's depiction of the situation is skewed either intentionally or unintentionally. And the purpose of limiting use of anonymous sources is to allow the public to evaluate the sources on which a journalist relies for a given story, seeing whether their perspective is likely to be

accurate and honest and whether an appropriately diverse range of perspectives was consulted in the formation of that story. Again, the value of these procedures, and of perspectival diversity in general, comes primarily from the same principle of chronic fallibility acknowledged by any convincing account of normative epistemology and emphasized by liberal political theorists for centuries. Because humans are both epistemically and morally fallible, we should not rely too heavily on any given grounds of belief or any given set of authorities. Instead, we must always preserve conditions wherein multiple perspectives can be brought to bear on any given issue so that errors made in good faith can be corrected and abuses of authority made in bad faith can be discovered and resisted.

As valuable as perspectival diversity is as a means of improving the reliability of our beliefs and discovering and resisting abuses of authority, however, insuperable difficulties brought on by the human condition limit the utility of perspectival diversity as a means of vetting testimony under conditions of weak consilience. Let us once again use the news as an illustration. It is obvious that in the present era much of the news media, particularly cable news, is influenced by partisan considerations. The likes of *MSNBC* and *Fox News* wear their partisanship on their sleeves, but many believe (myself included) that most other legacy news organizations tend to prioritize certain stories, adopt certain tones, and use certain language in a way that “leans” them away from full neutrality. One might suppose – and indeed I have often heard it argued – that the way for citizens immersed in such a world to get an “objective” view of the world is by consulting news sources with a variety of such biases and using a sort of process of epistemic triangulation to home in on the common truth that, it is presumed, lies somewhere in the middle. This approach would seem to be a way of bringing perspectival diversity to bear on a situation in order to improve the reliability of citizens' beliefs in the way I have just recommended, and indeed I think there is much to recommend it. However, I also believe it illustrates the limits of attempting to use perspectival diversity to ensure one adopts sound beliefs on the basis of testimony under conditions of weak consilience. One of the main problems is that *true* perspectival independence between sources is difficult to ascertain. This is not because all news is secretly dictated by some single, all-powerful source that secretly pulls the strings in a grand conspiracy to deceive the public and keep it under sugjugation (at least, not as far as I know). Instead, the problem arises from careful consideration of the news-production

process. For one thing, many of the reports disseminated by “independent” news organizations are sourced from the same original sources – the same interest groups, governmental authorities, police precincts, and so forth – as all the others. Indeed, the extent to which the press is dependent on a narrow and fairly uniform array of sources can be surprising to uninitiated. Consider Leif Wenar's description of the dilemma faced by those who consider donating to aid groups working in foreign companies. Such donors *appear* to have two separate sources of information about the effectiveness and practices of aid groups – the information provided by those aid groups themselves and the information provided by the media about those aid groups. But in reality the media – *all* the media – generally get their information *from* these same aid groups.⁴¹⁴ Why? It is a simple matter of economics. Even large media companies could not support the costs of independently conducting in-depth investigations of the practices of groups in areas as disparate and remote as Mongolia, Zambia, Estonia and northern Canada, so instead they make the practical, and perfectly understandable, economizing move of making relationships with a small number of readily recognizable and comparatively easy to access sources in each area and rely primarily on them for the information in their jurisdiction. Most often, such relationships are not formed by multiple independent media companies, but instead a single member of the Associated Press, from which the rest of the press receives the reports on which they base their published version of the stories they publish about that geographical area or group of interest.⁴¹⁵ This economically motivated reliance on such a narrow range of original sources is not, in my opinion, a case of moral failure or profiteering at the public's expense on the part of the media companies. Instead, it is simply a sign of how massive and complex the world the news tries to cover truly is – a massive and complex world in which no one can predict which events of public interest will happen, in which locations, and at what times. Even at the height of the legacy news era and running at a loss in terms of revenue, as has often been the case for news departments of major media companies, individual news organizations simply could not employ enough on-the-ground reporting teams to cover every single locale news conceivably wind up happening. This means they always have, and are for the foreseeable future

⁴¹⁴Wenar (2010)

⁴¹⁵This is more true than ever in our current era, which has seen a massive decrease in the scope of regional reporting departments among major legacy news organizations. However, it has always been the case to a large degree outside nations' capitals, political hot-spots and large metropolises. See Davies (2008).

always likely to, exploit the division of labor made available by use of the Associated Press and developed relationships with a comparatively narrow range of informants as described above. This in turn means that most of the news stories that are broadcast to the public are not gathered by the proximal news agency *independently* but which are, rather, simply *that agency's version of* a story provided both to them and to other news outlets by a narrow range of trusted informants and associates.

All these measures employed by the legacy news media are understandable, even admirable ways of dealing with the enormity, complexity, and unpredictability of the world they are asked to cover. It makes sense in such a world that major media organizations rely on built-up relationships with the Associated Press syndicates and local informants to cover the most ground they can in the most reliable way practicable. But as understandable and even commendable as it is, the fact still remains that mainstream news organizations seldom will or even *can* send their own investigative team to independently verify the reports they receive from such sources. Consequently, if a single source reports a story in a remote area is not true, or is grossly distorted, or is wildly misinterpreted due to cultural misunderstanding, we have no great assurance the mistake will ever be found out in the first place, let alone that the error's discovery will meet legacy news agencies' thresholds for "newsworthiness," especially when such a revelation may well undermine the perceived reliability of all the news organizations who disseminated the claim in the first place. Since the resources available to the the many news agencies that rely on a the story provided to all of them by a single AP syndicate or informant are limited for double-checking the details of that story, the veracity of *all* the news reports from *all* the agencies reporting that story is likely to depend heavily on the *same* factors, which means that any flaws in the original story-gathering process are therefore likely to be present in *all* versions published by *all* outlets. Thus, for much of the news, consulting "multiple sources" of news to get an "objective" sense of reality is something of a misleading characterization, since what the news consumer is actually doing is consulting *multiple versions* of a story produced by the *same source(s)*. Since different news agencies will have different editorial processes and thus couch their stories in different terms, rejecting and emphasizing different details, reading multiple agencies' versions of a given story is not *entirely* useless. However, this heavy

dependence on a single source or narrow range of sources does limit their ability to improve epistemic reliability by consulting multiple “independent” sources in the first place.

Similar considerations illustrate the difficulty of obtaining perspectival diversity by simply consulting multiple legacy news organizations, at least in any meaningful sense, in another way. In many areas, few outside the AP syndicates themselves and the informants they favor have (or *had*, before the internet era) a voice that is “loud” enough to catch the attention of the legacy news media. Every local AP syndicate or regular news informant is thus a locus of *power* that filters which information gets amplified and becomes available to large news organizations. As postcolonialists, sociologists and ethnographers have forcefully argued in recent decades, there is every reason to believe that there are many voices that are at best ignored and at worst actively stifled or distorted by these mechanisms of power.⁴¹⁶ By necessitating that these people's stories be told by interlocutors such as AP syndicates and major news organizations, the newsmaking process causes their diverse and direct on-the-ground experiences be filtered through the (unconscious and conscious) biases and agendas of both these local loci of power and of the news organizations themselves. Thus, what may well be quite a diverse array of unique occurrences, representing the experiences and perspectives of a broad range of individuals and groups, can by the news-making process itself condensed into a uniform set of narratives that fit into the worldview and/or agendas of a much smaller group of entities, many of whom share a common perspective due to their cultural upbringing, education and professional training. That is to say, newsmaking by the range of professionalized institutions that comprise the legacy news media is a *perspectivally homogenizing* process that diminishes the degree to which the information that passes through it may retain its capacity to bring perspectival diversity to bear on the news.

In the era that preceded the internet, relying on the information-gathering and disseminating processes used by the legacy news for the bulk of their information made sense, not because it was epistemically *ideal* but instead because humanity's technological means for gathering information and disseminating it to citizens required large, labor-intensive industries like that developed by the legacy news media. Given the resource-intensive nature of information gathering and dissemination under such conditions, only a few large companies could do it. In

⁴¹⁶For a powerful illustration of this, see Spivak (1988).

that world, it made sense for companies to rely primarily on internal quality controls and professionalized training to improve the reliability of the reports they eventually published, and citizens had few options apart from either trusting in those processes to get them reliable information or giving up on the idea of knowing anything about the world at all. Now, however, information dissemination is vastly cheaper and more accessible to entities beyond such huge and capital-intensive industries. While it is true that these decreased costs of dissemination mean that the information that makes its way to the eyes of citizens does not need to first pass through a rigorous process of quality control such as that valued by Goldman, it also means that citizens' ability to gain access to a diverse range of perspectives in the current era is *vastly* increased over that of the pre-internet age and that those perspectives need not go through the homogenizing process employed by the legacy news. This means that groups whose voices did not gain the ear of those who controlled the channels of information in the previous era have a chance of being heard in the present one and that democratic citizens have resources for forming a fuller, more reliable picture of the world, one less dependent on the judgment and good will of a small range of authorities than was the picture of the world formed by citizens in the past. It also means citizens have means of evaluating the performance of those authorities unavailable to the citizens of the past.

There is every reason to believe that this process works to improve the beliefs eventually formed by the public, both about the world at large and about the performance of the legacy news media. Take, for instance, a story that went viral in early 2019 about an interaction between a group of teenage boys wearing MAGA (“Make America Great Again,” Donald Trump's campaign slogan) and a Native American. The version of the story that originally went viral was based on a series of pictures and short video clips showing a group of high school students mocking a Native American man, Nathan Phillips, as he somberly played a drum and sang traditional Native American songs. *NPR's* original coverage of the story, titled “Video of Kentucky Students Mocking Native American Man Draws Outcry,” was typical of legacy news stories, relying heavily on Phillips' account of the incident and writing it in such a way as to encouraging readers to interpret it as evidence of racism and bigotry among Trump supporters, for instance mentioning that Phillips “told the *Post* he felt threatened and that the boys swarmed around him” and emphasizing that the students involved were “predominantly white,” wearing

“Trump paraphernalia” and were present in the area to attend the anti-abortion March for Life.⁴¹⁷ After the story, framed in this way, had been published by multiple legacy news outlets, attention was drawn to a separate video of the incident, which had been taken by a lay citizen who happened to be in the area for other reasons. That video showed that it was actually Phillips who initiated the interaction with the students, marching up to them without warning and playing his drum directly in their faces.⁴¹⁸ Moreover, the second video showed that a group who called themselves Hebrew Israelites had been preaching on the street and antagonizing both the students and the Native Americans for some time before the incident between Phillips and the students occurred. In the fuller context, the reactions of the 15 and 16 year-old students, who began jumping and chanting in time with Phillips' drumming and one of whom, Nick Sandmann, was caught smiling as Phillips sang in his face in a photo that went viral, seem at least as plausibly interpreted as teenagers reacting awkwardly and immaturely to a genuinely tense and bizarre moment than as signs of white supremacy and bigotry amongst Trump supporters, as they were originally interpreted in legacy news accounts. The release of the second video via Youtube caused a viral backlash against the original story, which itself caught the attention of the legacy news media and spawned the release of articles questioning the narrative implied by the original viral story, such as that released by the *New York Times*, titled “Fuller Picture Emerges of Viral Video of Native American Man and Catholic Students.”⁴¹⁹ This case illustrates the value of the sort of perspectival diversity enabled by the internet, both in terms of allowing citizens to gain a more nuanced view of the world and in terms of enabling citizens to assess the performance of the legacy news media. Without the release of the second video via Youtube, it is not at all clear that the narrative encouraged by the original news accounts would ever have been complicated or their rush to interpret it as an instance of bigotry by Trump supporters contested.

In order for perspectival diversity to have this kind of value, however, lay citizens need to be willing to question and contest the claims of the legacy news media and to consider with an open mind whether alternative accounts that appear via internet-era sources like Youtube and social media are a more accurate depiction of, or supplement to, the events reported by the news. This willingness is actively discouraged by the attitude of *selective skepticism* embraced by

417Paris (2019)

418 “Shar Yaqataz Banyamyan Facebook Video” (2019)

419Mervosh and Rueb (2019)

Progressives and defended by Goldman, which pre-defines certain privileged sources of information, such as the “scientific community” and the legacy news media, as “more reliable” and encourages the public to place its trust in them for the good of democracy and pre-defines other sources of information, such as social media and the blogosphere, as “unreliable sources” and encourages the public to adopt a disposition of skepticism and mistrust toward them. Instead of this kind of selective skepticism toward *some* kinds of testimony from *some* sources, citizens ought to subject *all* testimony from *any* single source to scrutiny in the light of their own judgment. One of the best tools they have for informing that judgment is to compare that testimony to testimony they believe to stem from a source or sources whose perspective is meaningfully independent of that of the original testifier. While the speed with which testimony travels makes it difficult to ascertain whether the other sources of information are *truly* independent from the original, and therefore lay citizens' confidence in *any* particular conclusion reached on the basis of such tools in the same way I have argued it ought to be tempered for all beliefs adopted under conditions of weak consilience, citizens' ability to make the most of perspectival diversity still depends critically on their willingness to subject testimonial claims to scrutiny in the light of their own judgment rather than adopting the attitude of excessive credulity toward favored epistemic authorities advocated by Goldman and Progressives.

Section Five: The Absence of Identifiable Expertise About News

We have looked at what I believe to be three of the most useful tools available to *anyone* attempting to adopt true beliefs on the basis of testimony under conditions where *strong* testimonial consilience cannot be had. These three tools are (1) consultation of a source's *reputation*, (2) vetting claims through one's *worldview*, and (3) seeking out *perspectival diversity* between testimonial sources. As I believe these are among the *best* tools available, use of whatever other tools there are is likely to be subject to the same or more profound limitations. Because these tools are *useful* and we live in a world where we must sometimes act even under non-ideal epistemic conditions, we should *not* suppose that placing one's trust in a given testimony or testifier on the basis of just one or a combination of these tools is always

unjustified. While *justified*, however, our *confidence* in the particular beliefs we adopt and the trust we place in testifiers and testimony chosen under such conditions ought to be *limited*.

Were this my whole argument, many readers would no doubt be unsatisfied. After all, doesn't the principle of chronic fallibility I have repeatedly emphasized imply that our confidence in *all* our beliefs ought to be limited, no matter whether they are formed under conditions of strong consilience, weak consilience, or any other variety of consilience a theorist might dream up? I ask that readers who feel such frustration creeping in indulge me just a little longer, because I believe once one further consideration is made the implications of my exegesis here will be both clear and weighty. This further consideration has to do with the role played by *judgment* in use of any of the tools described above.

So far I have emphasized the limitations inherent in attempting to use the three tools mentioned above – and by extension *any* tools available to us mere mortals – for forming beliefs on the basis of testimony under conditions where *strong* testimonial consilience cannot be obtained. There is one commonality between such scenarios, beyond these limitations, which, when considered in tandem with them, carries profound and far-reaching implications for contemporary politics. This is the fact that, no matter which of these tools is being used, individually or in combination, *no individual or group can be identified for which a persuasive case can be made that that individual or group's employment of these tools can be consistently relied upon to reach conclusions more reliable than those reached by the lay citizen*. That is, there is an *absence of identifiable expertise* when it comes to deciding who *really* deserves a reputation for trustworthiness, whose worldview is the *correct* worldview, or whether a given array of testifiers ought to count as sufficiently independent in the sense and to the degree required to justify the public at large's placing a high degree of confidence in their testimonies. No persuasive case can be made – that is, no case that relies on criteria whose merit is likely to be, or ought to be, universally agreed upon – that any identifiable individual or group is consistently better at making these sorts of judgments than lay citizens themselves. This makes immediate intuitive sense given the limitations described in the first part of this section. Since *everyone's* tools for evaluating testimony under conditions of weak consilience are subject to such substantial limitations, it makes sense that *no one* can lay claim to being an “expert” in that art, at least not in the sense that everyone else ought to feel compelled to take such would-be

experts' judgments about which testimonies to trust much more seriously than their own. I will subsidize this intuition with a more in-depth argument in a moment, but first let me provide a formal definition of the term “epistemic authority.”

Epistemic Authority: An individual or institution whose judgment in a given domain is so consistently and demonstrably more reliable than that of lay persons that the latter ought to routinely defer to the claims of that individual or institution, even when its claims contradict conclusions those lay persons would have reached had they relied solely on their own judgment

As I have defined the terms, to call someone an *expert* is to identify them as consistently having better judgment in a given epistemic domain than that of the lay person. Merely *having* better judgment, however, is not enough to be considered an *epistemic authority*. In addition to *having* better judgment than that of lay citizens, an epistemic authority's expertise must be *demonstrable to the public* by recourse to clear evidence. Otherwise, it violates the justified authority principle I identified as a fundamental element of liberal democracy, which states that in order to count as legitimate, all authority in a liberal democracy must be justified by reason *to* the citizens whose deference to that authority is demanded. By encouraging an attitude of routine acceptance of the claims of the legacy news media, the arguments of Goldman and Progressives imply that they believe the legacy news media to be *both* experts *and* epistemic authorities in the senses defined above. That view is suggested even more strongly by Progressives' growing support for censorship by “fact-checkers” of the information that is allowed to be exchanged online via social media platforms in order to combat “misinformation” and “fake news.” To suggest that some individual or group of individual ought to be appointed to decide for the lay public which news it is *allowed* to share and *allowed* to see is to suggest that that group's *judgment* about which news is true, and which sources of news are reliable, is not only superior to the public's but is so consistently and demonstrably superior to that of the public as to justify the public's routine deference to that group's decisions. Beyond this, it implies that Progressives believe that whatever “fact-checkers” are appointed to do this job both are presently and can be expected to remain indefinitely *trustworthy*, in that they will conduct their fact-checking duties in

good faith, even in the absence of oversight by the public. This must be the case because censorship by its very nature *precludes* oversight.

I believe Progressives err deeply in many ways by adopting this view. First, it is not clear whether there *are* experts about the news. Second, even if there are, I deny that their expertise can be *demonstrated* to the public in a way that even approaches the threshold required to meet the justified authority principle. Third, even if it could, censorship would *still* not be justified because censorship precludes oversight, and thus even in the unlikely event that any set of epistemic authorities could be trusted to do the job of censorship both honestly and effectively, they would inevitably become corrupt over time, which would be an enormous detriment to the well-being of democracy.

Let us take these arguments one by one. First, let's examine the notion of expertise about which news is true and which news sources are reliable. One notion of expertise has to do with people's mastery of a set of tools and techniques relevant the pursuit of knowledge in a given epistemic domain, along with the mastery of basic knowledge and concepts relevant to that domain. A legal expert, for example, does not know *all* the law – no one does – but they *do* know fundamental principles *of* the law as well as how to rapidly conduct research on the legal issues they are presented and connect that research to those fundamental principles. Likewise, a medical practitioner cannot know *everything* about human health, but they know many fundamentals of human health as well as the conditions common to different types of demographics, and they are able to conduct tests that will allow them to narrow down which ailments might be troubling their patients and direct them to the appropriate treatment(s). Neither legal scholars nor medical practitioners are perfect, but their judgments on legal and medical affairs generally surpasses that of the lay citizen to a great degree, largely because of their mastery of the basic knowledge and concepts relevant to their respective epistemic domains and their adeptness at using the tools and techniques relevant to the pursuit of knowledge in that domain. Expertise, then, would seem to consist largely in having mastered the basic knowledge and concepts relevant to a given epistemic domain, as well as the possession of a certain degree of adeptness in the use of the tools relevant to the pursuit of knowledge in that domain. When someone's mastery of the basic knowledge and concepts relevant to a certain epistemic domain and their adeptness in the use of the relevant tools for pursuing knowledge in that domain are

sufficiently great that their judgments in that domain come to be far more reliable than those of lay citizens, they count as an expert.

Should “fact-checkers” count as experts in this sense? It is not clear that they should. First, the epistemic domain of news is far more vast than the limited epistemic domains of the law and of human health. The epistemic domain of news includes *both* the law *and* human health, as well as the entire range of subjects that together comprise the realm of human affairs. It is doubtful whether anyone can be said to have “mastered” the basic knowledge and concepts relevant to the *entire realm of knowledge relevant to human affairs*. Second, when it comes to assessing the news, “fact-checkers” are limited to the same epistemic tools available to any of us under conditions of weak consilience. “Fact-checkers” have no more ability to personally inspect the events and phenomena reported in the news than any of the rest of us. Nor can they develop relationships with more than a bare fraction of the people who research, write, *or* edit the stories for a *single* news organization, let alone keep watch over *all* these processes, let alone do so for *more than one* organization so that they could meaningfully *compare* the reliability of their reports to those produced by others. Like the rest of us, “fact-checkers” are limited to the tools of weak testimonial consilience, which means that, like the rest of us, the consultation of a source's reputation, their own antecedent worldview, and the seeking out of perspectival diversity between testimonial sources are among “fact-checkers” best tools for deciding which news sources and stories are likely to be reliable and which are not. Can *anyone* be said to be significantly more adept in the use of these tools than lay citizens? It is not at all clear that they can. It is, therefore, not at all clear whether the idea of expertise about the news implied in the very idea of “fact-checkers” makes sense in the first place.

To say that it is “not at all clear” is, of course, not the same thing as saying that such expertise *does not* exist. Perhaps it does. Perhaps there really are some people who possess a superior understanding of the basic knowledge and concepts that underlie human affairs in general and have mastered the epistemic tools of weak consilience to such a degree that their judgments about which news is true and which news sources are reliable far outstrips that of lay citizens. Even if that is so, however, there is no guarantee that these will be the people who wind up being appointed as “fact-checkers” by Google, Facebook, Youtube, the government, the public, or any other appointing entity. This is because whoever (if anyone) possesses this kind of

expertise about the news cannot *clearly demonstrate* their expertise to the rest of us. The experts to which lay citizens tend most readily to defer in everyday life are those whose expertise is able to be made clearly manifest to the lay citizens doing the deferring. Medical practitioners' expertise is made clearly manifest to citizens when the practitioner's recommendations successfully resolve a medical issue the citizen was unable to resolve herself. Contractors' expertise in building homes is made clearly manifest to citizens in the form of houses they have built. Engineers' expertise in building bridges is made clearly manifest to citizens in the form of bridges, skyscrapers and countless other constructions most of us have no clue how to even begin building. The evidentiary clarity of these examples' expertise is provided by the ability of citizens to draw from their *personal experience* – to *feel* their symptoms resolved, to *walk through* a sturdy house, to *use* a bridge – in their assessments of the relevant entity's claim to expertise. Since the superiority of the judgments of these kinds of experts can be so clearly made manifest to practically any citizen, a stronger case can be made that citizens ought usually to defer to medical practitioners when it comes to physical health, contractors when it comes to building houses, and engineers when making bridges and skyscrapers. But the expertise of someone who claims to be especially skilled at determining which news organizations deserve a reputation for reliability cannot be made manifest to citizens with anything remotely resembling such clarity.

How can citizens *tell* if would-be “fact-checkers” judgments are more reliable than their own? Lay citizens have no better tools available to them for evaluating the performance of “fact-checkers” than the tools of weak consilience mentioned above – consultation of the would-be “fact-checker's” reputation, checking whether the “fact-checker's” claims make sense in light of the citizen's worldview, and examining the degree to which the “fact-checker's” claims are compatible with the claims of other sources the citizen trusts. That citizens must rely on such tools is a problem for two reasons. First, while lay citizens are not *powerless* to form reliable beliefs under conditions of weak consilience, I deny that the epistemic tools of weak consilience generally meet the standard of evidentiary clarity required to meet the standard of public justification. Having one's ailment cured after taking the advice of a physician is a much stronger evidentiary basis for considering that physician to be an expert than simply being familiar with the reputation of a given newspaper or “fact-checker.” Walking through a house built by a given

contractor or having a friend who lives in a house built by that contractor are stronger evidentiary bases for accepting that contractor as an expert in home construction than a news source's or "fact-checker's" tendency to issue reports in a way consistent with one's antecedent worldview. Crossing bridges and seeing skyscrapers regularly fail to fall down are much stronger evidentiary bases for belief that the engineers who design those structures deserve to be considered experts in their fields than simply knowing that multiple news sources' stories about a given event or phenomenon match, or that multiple "fact-checkers" one supposes to have nothing to do with one another agree that a given news story is true or a news source reliable. In each of these examples, the evidentiary clarity available to lay citizens through *personal experience* is generally sufficient to meet the standard of public justification for treatment of the relevant experts as epistemic authorities, while the evidentiary clarity available to lay citizens through the *epistemic tools of weak consilience* is not.

The second problem with the idea that "fact-checkers" expertise is sufficiently demonstrable to lay citizens to meet the demands of public justification is that, just like I argued above in objecting to Progressives who view the legacy news media's reputation for truth-telling as a good reason to trust its claims while simultaneously viewing the judgment of lay citizens as deeply unreliable, it implies self-contradictory views about the judgment of lay citizens. Figuring out who *really* deserves to be treated as an expert when it comes to "fact-checking" depends on citizens making good judgments about which testimonies they ought to trust under conditions of weak consilience. But the *entire basis* for appointing "fact-checkers" to censor citizens' information in the first place is the supposition that lay citizens' *own* judgment under conditions of weak consilience is unreliable, otherwise they should be able to decide which news is true and which news sources are reliable for themselves. If justifications *of* the epistemic authority of "fact-checkers" must appeal to the very sorts of judgment in which lay citizens are supposed to be deficient, then, there would appear to be an irresolvable dilemma. Either lay citizens *are* fit to assess the reliability of "fact-checkers," in which case there would seem to be no need for "fact-checkers" to decide which news is true and which news sources are reliable in the public's stead, or the public is *not* fit to assess the reliability of "fact-checkers," in which case the "fact-checkers" authority cannot be justified to the public by recourse to reason and therefore fails to meet the demand of public justification. Not only, then, do I believe we have good reason to be

skeptical of the very idea *that* any particular individual or group actually *is* an expert about which news is true and which news sources are reliable, but even if such experts *do* exist, their expertise is not *demonstrable* to the public in any way that meets the demand of public justification, and therefore such an expert or group of experts should not be treated as an *epistemic authority* in the domain of news. Therefore, there is no compelling reason for the public to defer to the judgment of anyone else who has an opinion about which sources of information ought to count as “reliable” rather than relying on their own, and there is *certainly* no justification for the imposition of “fact-checkers” as censors of the information available to the public.

Beyond this, even if we *did* believe there to be experts about which news is true and which news sources reliable and *could* justify their being treated as epistemic authorities, we would *still* have compelling reasons to oppose both censorship *by* those epistemic authorities and the adoption by the public of an attitude of routine deference to them. The situations in which expertise is *most* useful are precisely those situations in which our *own* judgment is *least* useful, from which it follows that the importance of relying on the judgments of experts derives entirely from those situations in which our own judgment *differs from* the recommendations of the experts. That is, all situations in which expertise is of any practical use to the lay public are situations in the public decides to *forfeit its own judgment in favor of the judgment of the expert(s)*. All situations in which such a forfeiture is entailed carry the risk of sanctioning decisions we ourselves would not choose if we had full knowledge, and all carry the risk of enshrining authorities who may take advantage of our trust to further their own plans. As I have already argued, these risks are only amplified if lay citizens demonstrate a consistent tendency to routinely defer to epistemic authorities, for such a tendency must inevitably be noticed by those authorities and, once it is, holds out such an obvious temptation for those authorities to abuse their power that we must surely expect it them to do so sooner or later. The situation is even worse if those authorities are allowed to censor the information available to citizens, in which case the temptations held out to the authorities would be even greater and the consequences even worse, since citizens would not have the resources required to conduct oversight of those authorities even if they wished to do so. Even if we suppose that there are experts in the epistemic domain of news and that that expertise could be demonstrated to the public clearly

enough to meet the demand of public justification, then, citizens *still* have compelling reasons to oppose the adoption of the attitude of routine deference to “fact-checkers” of the sort favored by Progressives and, *a fortiori*, to oppose granting those “fact-checkers” the authority to censor the information available to the public in the way Progressives have increasingly come to support.

Conclusion: Trust, Testimony and Authority in Liberal Democracies

From the arguments outlined in this chapter I draw the following conclusions. First, *deciding which news stories and news sources to trust, in which situations, and why is a matter on which there is no identifiable expertise*. Since no individual or group can be identified for which a compelling case can be made that they are bound to be better at figuring out which testimony or testifiers deserve to be trusted, it follows that, second, *decisions about which news stories and news sources to trust, in which situations, and why ought to be left to the judgments of citizens*. To do otherwise excessively enshrines authority and cultivates habits of deference to entities whose authority cannot be publicly justified in ways that are incompatible with the aims of liberalism and democracy.

This in turn means that *no filtering of information by authorities on the basis of fact*⁴²⁰ *can be justified in a way consistent with the ideals of liberal democracy*. Finally, it means that *lay citizens have no compelling reason to suppose that anyone's opinion about which sources of testimony to trust ought to be given more prima facie weight than their own*. None of this, of course, implies that no one can be *better* than anyone else at deciding what testimony to trust. It just means we ought to expect people's opinions about who really *is* better at this task to differ, and that we have no *prima facie* reason to favor one side of the argument over the other. Instead, such disagreements are what John Rawls famously called “burdens of judgment,”⁴²¹ that is, disagreements between reasonable persons that cannot be adjudicated by appeal to any readily identifiable set of evaluative standards we ought to consider decisive.

⁴²⁰I say “on the basis of fact” to avoid getting into the topic of what is often called “hate speech.” My own opinion is that this should not be filtered by authorities either, for numerous reasons, but it is not necessary to delve into the topic in order to address the issues I am concerned with here.

⁴²¹Rawls (2005), Lecture 2, Section 2

In short, we have a host of reasons, both theoretical and practical, to embrace the following principles:

No Identifiable Experts About News Principle: There are no identifiable experts in the skill of determining which testimonial claims merit confidence across all epistemic domains encompassed by the news.

Trust as a Province of the People Principle: Decisions about which testimony or testifiers to trust ought to be left to the judgment of lay citizens.

No Prior Restraint Principle: No filtering of information on the basis of fact by any set of would-be authorities can be justified in a way consistent with the ideals of liberal democracy.

Throughout this dissertation I have emphasized the centrality to liberal democracy of a certain modest optimism about the judgment of lay citizens and the strong preference for doxastic self-rule while objecting to Progressives' deep pessimism about lay citizens' judgment and the sort of routine deference to certain established authorities implicit in their views of responsible citizenship. In this chapter, I have extended my theoretical defense of these positions and tried to clarify them through various illustrations. I have argued that, while the tools available for citizens under the conditions of weak consilience that characterize the vast majority of their decisions about which news stories and sources to trust are *imperfect*, they are still *useful* enough to justify a certain degree of modest optimism that citizens possessed of such tools will be able to use them to navigate the informational environment of the 21st century well enough to form largely reliable beliefs. At the same time, I have argued that the fundamental limitations faced by citizens under such epistemic conditions are the same that must be faced by *anyone* trying to decide which news is true and which news sources are reliable. I have argued that this ought to cast doubt on the very idea of expertise when it comes to the news, that such expertise cannot be demonstrated to the public in a way to justify such experts' being treated as epistemic authorities in the way suggested by Progressives, and that even if it could, citizens would *still* have compelling reasons to oppose the uncritical attitude toward those authorities demonstrated and advocated by Progressives, and oppose even more adamantly the placement of those

authorities into positions of censorship. Instead of either deferring as a matter of course to the opinions of some set of “fact-checkers” about which news is reliable or relying on those “fact-checkers” to exercise prior restraint on the information that becomes available to them online, lay citizens ought to be encouraged to put the epistemic tools of weak consilience to use and come to their own conclusions about which news stories are true and which sources reliable. Moreover, when the conclusions those citizens reach on the basis of those tools conflict with that of some set of “fact-checkers” or the depiction of that story by some source of news, citizens ought to feel free to object to the conclusions of those would-be authorities on the basis of their own judgment and not feel overly bothered by whatever accusations of epistemic malfeasance might be thrown upon them by fellow citizens who happen to have come to different conclusions on those same matters. Contrary to the belief of Progressives, who routinely throw out such accusations and charge lay citizens who decline to defer to their preferred epistemic authorities as *threats* to democracy, the willingness of citizens to exercise doxastic self-rule and object, on the basis of their own judgment, to the claims of would-be authorities is *integral to* democracy's well-being.

In the next, and final, chapter, I will argue that the problems characteristic of the Progressive position stem largely from their having adopted an epistemic outlook that underestimates the problem of chronic fallibilism. I will argue that by adopting an alternative outlook, one I call *epistemic pluralism*, we can avoid the problems of Progressivism and approach public policy and civic epistemology in a way that is at once more consistent with the ideals of liberal democracy and better poised to take advantage of the considerable potential for 21st century communications technologies to help us realize those ideals.

Chapter Six: Perspectivism and 21st Century Citizenship

“[W]hat is the faith of democracy in the role of consultation, of conference, of persuasion, of discussion, in formation of public opinion, which in the long run is self-corrective, except faith in the capacity of the intelligence of the common man to respond with commonsense to the free play of facts and ideas which are secured by effective guarantees of free inquiry, free assembly and free communication? I am willing to leave to upholders of totalitarian states of the right and the left the view that faith in the capacities of intelligence is utopian. For the faith is so deeply embedded in the methods which are intrinsic to democracy that when a professed democrat denies the faith he convicts himself of treachery to his profession.”

*John Dewey, 1939*⁴²²

In the first substantive chapter of this dissertation I argued that any convincing defense of liberal democracy must (a) hold a certain *modest optimism* about the judgment of lay citizens and (b) envision those citizens as exercising *doxastic self-rule*, that is, deciding what to believe and whom to trust on the basis of their *own judgment*. The argument of that chapter defended the connection between liberal democracy, modest optimism about lay citizens' judgment and the endorsement of doxastic self-rule on grounds of internal consistency. That is, I argued that liberal democracy's combination of protections for freedoms of communication and universal, equal opportunity of suffrage make no sense unless one envisions citizens as possessing reasonably sound judgment and being assertive in its exercise. I hope by now to have supplemented this argument centered on the *internal* consistency of liberal democracy with a compelling case that a liberal democracy built on this same modest optimism about lay judgment and the

⁴²²Dewey (1976 [1939]), 226

encouragement of doxastic self-rule is worth putting into practice in the *external* world. While I have defended a modestly optimistic view of lay judgment on empirical and historical grounds and repeatedly emphasized the value of encouraging citizens to assertively exercise doxastic grounds as a means of improving knowledge and discovering and resisting illegitimate authority, these arguments have appeared in disconnected bits and pieces. In this final chapter, I will collect what I take to be the most important of those bits and pieces and construct a different model of liberal democracy, and the citizenship appropriate to it, than the one that underlies the currently dominant Progressive model.

Section One: Chronic Fallibility

Fallibility as a Recurring Theme

A recurring theme in this dissertation has been the idea of *fallibility*. As a reminder, when I use the term “fallibility” by itself I mean it to signify *both* epistemic fallibility – that any one of our beliefs might be wrong – and moral fallibility – that human beings are often lazy, apathetic, corrupt, or beset by any combination of other vices. I dealt with the concept of fallibility most explicitly in Chapter Five, where I discussed the implications of *chronic* fallibility and used those thoughts to derive a normative model of belief formation wherein the *amount of confidence* an agent is justified in attaching to a given belief ought to be based on the *degree of consilience* available to that agent in support of that belief. However, the concept of fallibility has been crucial in every other chapter as well. Obviously, the fallibility of lay citizens is at the forefront of Progressives' worries about the plight of 21st century democracies and heavily implicit in the extensive social science literature on cognitive biases, which I have mentioned in several chapters. The concept of fallibility is also at the forefront of my account in Chapter Three of gastroenterologists' research on peptic ulcer disease (PUD) and the characterization of the human quest for knowledge as a “bandit problem” with unknown payouts for different avenues of research, as well as my arguments about the best ways for us as a society to improve our collective knowledge and discover illegitimate authority. There is a key difference, though, between the sort of fallibility emphasized by Progressives and the sort of fallibility I have emphasized. The fallibility emphasized by Progressives is the fallibility of *lay citizens*, whereas

the fallibility I have emphasized is the fallibility of *all human beings and human-created institutions*. It is worth dwelling a short while on this difference in emphasis, as it bears vital implications both for public policy as well as for civic epistemology.

Institutionalism in Progressive Thought

In Chapter Two I described what I called the “Progressive Ideal of Democratic Epistemic Well-Being,” defined as follows:

Progressive Ideal of Democratic Epistemic Well-Being: A state of affairs in which citizens have access to and generally act on the basis of reliable information, which importantly includes a propensity to consult and believe expert and/or professional knowledge-dissemination institutions.

In that chapter I first argued that Progressives heavily emphasize the importance of knowledge-producing and disseminating institutions, and the public's willingness to place their trust in them, to the well-being of democracy. After giving further evidence of this proposition in each subsequent chapter, I hope the point has been sufficiently demonstrated. I call this facet of the Progressive outlook, which emphasizes the importance of such institutions in maintaining the epistemic well-being of democracy, *institutionalism*.

Progressives' esteem of their favored institutions is so strong, in fact that many seem to equate the refusal to defer to those institutions' claims as equivalent to denying the very idea of objective reality. This seems to be the sort of thing, for example, that Stephan Lewandowsky and colleagues are worried about when they warn that one of the dangers of the post-truth “malaise” is that citizens are beginning to “stop believing in facts altogether.”⁴²³ It also seems to be the source of the sensational reaction to the statement by then-President Trump's adviser Kelly Anne Conway that former Press Secretary Sean Spicer's wildly inaccurate claims about the number of people who had attended Trump's inauguration were just “alternative facts,” which for many Progressives encapsulated all that is wrong with the post-truth movement.⁴²⁴ Thus, Progressives

⁴²³Lewandowsky et al (2017), 354

⁴²⁴The reader can ascertain this for herself through a quick search of the term “post truth and alternative facts” on Google Scholar, which will reveal dozens of books and articles published in the scant four years since 2017 with both those terms in its title.

tend to see the only alternative to deferring to some set of institutional authorities as *epistemic relativism*, the notion that *everyone's* beliefs about the world are *equally* valid. This exact sort of worry seems to be straightforwardly what drives Tom Nichols to warn about the dangerous growth “of an irrational conviction among Americans that everyone is as smart as everyone else”⁴²⁵ and Peter Dahlgren's condemnation of climate deniers, whom he views as paradigmatic examples of the post-truth condition, for viewing “the overwhelming scientific evidence” as just “another subjective opinion.”⁴²⁶ Progressives, then, tend to see institutions as the primary determinants of objective reality and the refusal to defer to their favored institutions as a descent into epistemic relativism.

Now, I do not *entirely* disagree with Progressives about the value of institutions in helping citizens form an accurate view of the world. Like Progressives, I believe that in many, if not all, epistemic domains, *some* methods are bound to be better than others at producing reliable knowledge about the real world and that there can be massive benefits to forming institutions dedicated to developing and employing such methods to seek out knowledge in those domains. I also grant that the scientific community and (the best examples of) legacy journalism are in fact two such institutions. In addition, I agree with Progressives that epistemic relativism is a problematic epistemic framework and would be dangerous to democracy were it to be embraced by large numbers of citizens. Like Progressives, I believe some propositions about the world are more accurate than others.⁴²⁷ Likewise, some institutions' claims about the world are consistently more reliable than others. For a large bulk of citizens to deny this would not only be self-deluding but disastrous, as it would wreck their ability as individuals to improve their beliefs and hamper our collective ability to correctly identify and form reasonable solutions for societal problems. Given the size and complexity of the world, institutions that specialize in particular

⁴²⁵Nichols (2017), 7

⁴²⁶Dahlgren (2018), 25

⁴²⁷I personally prefer the term “useful” to the term “accurate” because I am sympathetic to the pragmatic conception of truth, which holds that a thing is “true” if it is a consistently useful basis of action, i.e. if it acting on its basis consistently helps us meet our desired goals. This is technically a different conception of truth than the conception of “correspondence to reality” embraced by realists and perhaps most intuitively implied by the term “accuracy.” However, I do not think this distinction is vital here, since my point is saying that “some propositions about the world are more accurate than others” is to insist that propositions about the world can often be qualitatively ranked. That claim makes sense no matter whether one views the standard by which they are ranked is in terms of their fidelity to reality or in terms of their ability to help us take effective action.

epistemic domains and develop methods tailored to effectively exploring those domains are indispensable.

My disagreement with Progressives is *not* that they view institutions as societally valuable – even indispensable – sources of knowledge about the real world, nor that they view epistemic relativism as both descriptively in error and potentially dangerous to democracy. Instead, my disagreements with Progressives are that (a) their *version* of institutional realism is not grounded in a sufficient appreciation of the problem of *chronic fallibility* and that (b) in many cases, behavior they interpret as evidence of a societal descent into epistemic *relativism* is better interpreted as the rise of a different, and beneficial, epistemic outlook, which I call *epistemic pluralism*. I propose that both these deficiencies could be shored up by adopting an approach to civic epistemology exemplified by two of the American founding fathers, James Madison and Thomas Jefferson, both of whose work placed a central emphasis on the problem of chronic fallibility and not only recognized but embodied and embraced the sort of epistemic pluralism miscategorized and condemned as relativism by Progressives. By contrasting the epistemic outlook of Madison and Jefferson with that of Progressives, I hope not only to demonstrate how the former avoids the problems I have identified in the latter, but can also serve as the basis for a very different – and in my view superior – interpretation of the epistemic impact of 21st century communications technologies on democracy and the policies we ought to put in place to best make use of those technologies' potential for improving the epistemic, and general, well-being of democracy.

Progressives' Selective Notion of Fallibility

As we have seen at numerous points in this dissertation, Progressives view 21st century communications technologies as posing a threat to the epistemic well-being of democracy largely because they perceive those technologies as undermining the public's trust in certain established epistemic authorities. Their treatment of this issue strongly implies that the sort of fallibility we ought to be worried about lies preponderantly on the *lay citizens'* side of the expert-lay citizen relationship. As much is implied in the many cases in which a commentator focuses heavily on the epistemic flaws of citizens and points out the importance of trust in established authorities without making any mention whatsoever of the possibility that those authorities might

be corrupt, in error, or both.⁴²⁸ But even when Progressives *do* buck this general trend and acknowledge the moral and epistemic fallibility of experts, they tend to (a) preponderantly emphasize the threat posed to democracy by the epistemic flaws of ordinary citizens than by the fallibility of epistemic authorities, (b) downplay the democratic danger posed by corruption and/or good-faith error among expert groups, and (c) recommend it be combated primarily through what amounts to self-policing. In the end, the message is still the same: lay citizens are depicted as possessing insufficient powers of judgment to make reliable assessments of the performance of experts and because of this are expected as a matter of civic duty to routinely defer to those experts.

Tom Nichols' book *The Death of Expertise* is an excellent example. Nichols does a better job than many Progressives of acknowledging the fallibility of his favored epistemic authorities and exploring the potential problems that fallibility might pose to democracy, dedicating a whole chapter to the topic. In that chapter, Nichols admits that experts can make good-faith errors,⁴²⁹ make claims that lie beyond their epistemic jurisdiction,⁴³⁰ and engage in “outright deception and malfeasance.”⁴³¹ Indeed, Nichols goes so far as to insist that there is “no way around” the fact that “a non-negligible amount of published scientific research is shaky at best and falsified at worst,”⁴³² and since his examples of the consequences of these various forms of fallibility on the part of experts include Americans' accelerated obesity and death as a result of bad health advice from nutritionists,⁴³³ one can infer that Nichols acknowledges that experts' fallibility can cause significant public harms. While these acknowledgments demonstrate that Nichols is not fully dismissive of the potential for experts to become corrupt or complacent, however, Nichols' thoughts on what we ought to *do* in light of experts' fallibility still leave very little room for lay citizens to exercise meaningful oversight over those experts in an effort to discover when those experts might be in error, or corrupt, or both. On the contrary, Nichols consistently ends up overlooking, downplaying and even condemning the efforts of lay citizens to subject experts' claims to critical scrutiny, charitably writing off experts' errors as part of the normal human

428E.g. Lewandowsky et al (2017), Schiffrin (2017), Dahlgren (2018)

429Nichols (2017), 175

430Ibid., 177-8

431Ibid., 178

432Ibid., 182

433Ibid., 171-2

learning process, and even backhandedly praising the same communities of experts whose failures he uses as illustrations. Take, for example, the anecdotes Nichols uses to illustrate the potential danger of expert fallibility and/or corruption in the chapter titled “When the Experts Are Wrong.” In the opening anecdote, Nichols tells of an eight-grade student successfully using “research she did on Google” to rebut the claim of a “distinguished historian” that the “No Irish Need Apply” signs commonly supposed to have been used by employers in the late-19th century had never actually existed. The eight-grader in question, Rebecca Fried, found her evidence simply by “looking through databases of old newspapers,” which neither the original researcher nor any other scholarly historian had “apparently bothered to do” in the 10-plus years since the original claim denying the existence of the anti-Irish signs was published.⁴³⁴ Another of Nichols' anecdotes tells of the debunking of an Emory historian's claims about the history of gun ownership in America, whose dubious use of sources might otherwise “have gone unnoticed...attracted closer scrutiny” because they aroused the attention of “gun control advocates and gun ownership groups.”⁴³⁵

One plausible moral to draw from both these anecdotes would be that in at least some cases vigilance on the part of individual lay citizens and interest groups composed largely of them can serve to make sure the experts are doing their due diligence and incentivize them to do so. Another plausible implication of the first anecdote might be that the internet has the potential to facilitate just this sort of watch-keeping on the part of citizens in unprecedented ways – for example, by giving eight-graders ready access to the archives of a wide variety of newspapers published more than a century ago. Still another might be that the judgment of lay citizens – even eight-graders – can at least in some cases be good enough to serve as a legitimate basis for those citizens to contest the claims of epistemic authorities. Nichols, however, draws neither of these conclusions. Instead, the sole purpose of these anecdotes, as well as other anecdotes in the chapter, is apparently to do nothing more than establish the rather bland point *that* experts *sometimes* make mistakes. At no point in the chapter does Nichols suggest that such failures on the part of experts, which Nichols himself says to be common, might suggest a legitimate role for *lay citizens* in maintaining vigilance over such communities, despite that being precisely the

434Ibid., 170-1

435Ibid., 186

means by which historians' claims about the absence of "No Irish Need Apply" signs in 19th century America were debunked in the anecdote that begins the chapter. Instead of concluding that there might be a legitimate role for lay citizens to play in watching over experts and/or that 21st century communications technologies might contribute to citizens' ability to make sure experts' claims are indeed reliable, Nichols immediately follows up his admission that experts "get things wrong all the time" by rushing to insist that nevertheless "laypeople have no choice but to trust experts"⁴³⁶ and goes on to spend most of the rest of the chapter backhandedly defending the mistakes of the experts in question while criticizing lay citizens in various ways. For example, Nichols' first response to the apparently alarming failure of large-scale efforts to replicate studies in the field of psychology is to pre-empt accusations of research fraud⁴³⁷ and cite a study by another group of social scientists claiming the replication study was "completely unfair – and even irresponsible."⁴³⁸ Only after having done this does Nichols briefly mention, in an offhanded way, that expert research is beset by other problems, such as "poor oversight of grants" and "intense pressure from academic institutions to come up with publishable results."⁴³⁹ It is telling that Nichols makes no suggestion that these structural problems might be the basis of legitimate doubts about the reliability of the very institutions to which he is asking the public to defer, or that oversight of these institutions by the lay public might be part of the solution. In fact, he draws no inferences about either the institutions in question or democratic citizens' relationship to them whatsoever. Instead, he merely mentions that the above-cited issues *are* problems before quickly moving on to other topics.

It is also telling that he finds space in the same section to squeeze in praise for the expert community whose research failed to replicate, noting that the "whole business" of discussing the original replication study, social scientists' rebuttal, and even "a rebuttal to the rebuttal, is now *where it belongs*: in the pages of the journal *Science*."⁴⁴⁰ I take the implication of the italicized phrase to be that contesting the claims of experts is the rightful province of other experts, and not of the lay public. The moral of the story, Nichols would have us believe, is that it is not clear whether the experts have erred in the first place, and even if they did, it is still their job to

⁴³⁶Ibid., 174

⁴³⁷Ibid., 184-5

⁴³⁸Ibid., 184. This quote is from Harvard political scientist Gary King.

⁴³⁹Ibid., 185

⁴⁴⁰Ibid., 184, my emphasis

critique their *own* work and publicize their *own* errors. The pages of the journals written and edited by the experts themselves – and *not* in the pages of blogs or other publications written and disseminated online by lay citizens – is where such criticism “belongs.” This point is reinforced at the end of the chapter, which concludes with a section that purports to be a discussion of the responsibilities of *both* experts and lay citizens when it comes to experts' fallibility. In reality, Nichols relays the responsibilities of experts in a single sentence and spends the rest of the section warning of the dangers of citizens trying to judge the claims of experts themselves and admonishing them to stay in their lane. The single-sentence summary of the responsibility of experts is that they must “own their mistakes, air them publicly, and show the steps they are taking to correct them.”⁴⁴¹ Meanwhile, laypeople “must exercise more caution in asking experts to prognosticate, and they must educate themselves about the difference between failure and fraud.”⁴⁴² They must also approach expert advice “with a certain combination of skepticism and humility,”⁴⁴³ and look to public intellectuals to serve as mediators and interpreters between themselves and the experts rather than either attempting to make judgments about the experts' claims for themselves or casting about for information on the internet, which Nichols characterizes as populated largely by “amateurs, hucksters, charlatans, and conspiracy theorists.”⁴⁴⁴ In other words, Nichols' recommendation for how to deal with the problem of expert fallibility and/or corruption reduces to offering a brief moral exhortation to experts followed by a plea to lay citizens to avoid relying on their own – highly suspect – judgment to assess the claims of experts, as well as a paternalistic admonishment for citizens to “educate themselves” in order to avoid making elementary category errors.

The point to note in all this is that even for Nichols, who is exceptional among Progressives in dedicating an entire chapter of his book to the problem of expert fallibility and/or corruption, the overall approach to the topic of fallibility is the same as that of other Progressives who make no such effort. By far the sort of fallibility of most concern to democracy is the fallibility of lay citizens, whose judgment is inadequate to the task of keeping watch over expert communities and is indeed so poor as to make the very act of seeking out information online a

441Ibid., 205

442Ibid., 206

443Ibid., 207

444Ibid., 206

doubtful prospect.⁴⁴⁵ Meanwhile, though Nichols makes some token gestures toward admitting that the fallibility and/or corruptibility of expert communities might be problematic, his proposed “solutions,” which consist entirely in a mix of moral exhortations for those communities to police themselves better joined by warnings to citizens of the dangers of *their* trying to exercise oversight over such communities, show that he does not really take the problem seriously. In other words, even though Nichols' *The Death of Expertise* is an exception to the general Progressive tendency to ignore the fallibility and corruptibility of experts, it is an exception that proves the rule that they do not take these possibilities or their potential consequences to democracy seriously.

In Chapter Four, I argued that Progressives advocated for citizens to adopt a disposition of *selective skepticism*, subjecting only information from *certain* sources, such as social media, to strict scrutiny in the light of their own judgment while approaching information disseminating from *other* sources, like the “scientific community” and legacy journalists, with an *a priori* disposition of credulity. Nichols' *The Death of Expertise* illustrates how this disposition of selective skepticism is paired with a similarly selective view of fallibility. For Progressives, the sort of fallibility that poses by far the biggest problems for democracy is the fallibility of lay citizens, which stems largely from their impoverished, or at least deeply unreliable, powers of judgment. The fallibility of Progressives' favored epistemic authorities, on the other hand, is consistently downplayed and addressed by non-solutions that reduce to empty moral admonishments and self-policing – if indeed it is even acknowledged at all. I call this tendency to emphasize the fallibility of lay citizens and to ignore, downplay or overlook the fallibility of experts and favored institutional authorities *selective fallibilism*.

Problems with Selective Fallibilism

The problem with this one-sided approach to fallibility is that it underestimates *both* the fallibility of epistemic authorities *and* the capacity of lay citizens to use their own judgment to identify good-faith errors and corruption among those authorities. That it underestimates the fallibility of epistemic authorities is evident both in the way Progressives tend to ignore and

⁴⁴⁵As with many other Progressives, the latter message is replete in Nichols' book, whose fourth chapter's subtitle reads “How Unlimited Information is Making us Dumber” (Ibid., 105).

downplay the very idea of fallibility among their favored epistemic authorities and, even in the rare cases when they *do* attempt to acknowledge it, wind up responding by offering empty moral exhortations and self-policing as “solutions.” No one who recommends self-policing as the solution to the possibility an institution might be corrupt, complacent, or possessed of epistemic blind spots can take that possibility very seriously. This is because, whether it be corruption, complacency, or the possession of an epistemic blind spot, the very condition self-policing is meant to address would preclude its being effective. To suggest self-policing as the solution to such a problem is to suggest that the problem either does not really exist or is not very serious. Yet, I take it as obvious that such possibilities *do* exist and that they *are* very serious when they occur. To consistently ignore, overlook and downplay them and to offer self-policing as the primary remedy for them on the rare occasion they are acknowledged is evidence that Progressives underestimate the likelihood and consequences to the public of the fallibility of their favored authorities. At the same time as it places excessive confidence in the moral and epistemic reliability of these authorities, the Progressive view is excessively pessimistic about the judgmental capacity of lay citizens, especially when it comes to their capacity to exercise meaningful oversight of the authorities favored by Progressives. This can be seen in the way they seem to fully ignore the numerous available examples of novices contributing to specialized bodies of knowledge, as in the “No Irish Need Apply” example above, and/or discovering corruption in expert communities, as was the case with the discovery that the *Lancet* article posing a causal connection between the MMR vaccine and autism was based on fraudulent data, as described in Chapter Three. It is as a result of this combination of underestimating the possibility and consequences of experts' fallibility and underrating citizens' judgmental capacity when evaluating the work of these authorities that Progressives like Nichols tend to settle for toothless oversight measures and admonish citizens to “stay in their lane” and not try to make judgments about matters on which they are not experts. The factor that drives all these tendencies is that Progressives' view of fallibility is *selective*. It is so rampant among lay citizens as to discredit their judgments over vast swaths of epistemic territory, yet so minimal among experts as to require few, if any, measures of mitigation.

In treating fallibility as *selective* in this way, Progressives depart from a deeper notion of fallibility that is at once a better portrayal of the human condition and a better guide for public

policy and civic epistemology. This is the notion of *chronic* fallibility, a notion that was central to the writings of James Madison and Thomas Jefferson. Far from being the unique province of the *hoi polloi*, for Madison and Jefferson both moral and epistemic fallibility were a fundamental and insuperable aspects of the human condition for *all* people, as well as for any institution created by and composed of them. Because they viewed fallibility as both insuperable and pervasive, and not something that could be reliably overcome by any amount of specialized training either in intellect or in ethics, these thinkers advocated approaches to public policy and civic epistemology that guarded against the dangers of complacency and corruption among established epistemic authorities in ways the approach of Progressives does not. The centrality of *all* human beings' epistemic and moral fallibility to these figures' thinking impelled them to resist *any* situation where a given set of authorities served as the sole or even primary means of its own oversight and to view moral exhortations directed toward those authorities as deeply inadequate to the task of resisting corruption and complacency among them. It also underwrote their uncompromising emphasis on freedoms of communication, even though they were far from holding romantically optimistic views of the judgmental capacity of the average democratic citizen.

The recognition that the epistemic and moral fallibility of human beings is *chronic*, neither limited to some particular class nor eradicable by any reliable means, drove Madison and Jefferson to address it in the same way: by stressing the need to foster *perspectival diversity*, both in decision-making bodies and in society as a whole. As I will argue, these thinkers' turn to perspectival diversity as the proper response to the problem of chronic fallibility reflected their shared embrace of an epistemic outlook, which I call *epistemic pluralism*, that allowed them to account for the fallibility of even the (currently) most reliable institutions without collapsing into the void of epistemic relativism feared by Progressives. The epistemic pluralism embraced by Madison and Jefferson embodied a *modest optimism* about lay citizens' judgment and encouraged them to assertively exercise *doxastic self-rule* in the ways I have named as central to convincing defenses of liberal democracy in this dissertation due to their contribution to the growth of knowledge and the discovery of illegitimate authority. As such, I believe the epistemic pluralism of Madison and Jefferson to be a better guide for public policy and our models of civic epistemology than the institutionalism of Progressives.

Diversity as the Proper Response to Chronic Fallibility in Federalist No. 10

Madison's *Federalist No. 10* is among the most widely read and cited of the *Federalist Papers*. In most cases, the main focus of the analysis of *No. 10* is either on Madison's warnings of the dangers of faction⁴⁴⁶ or his apparently unflinching defense of individual liberty despite those dangers.⁴⁴⁷ There are, of course, many other scholarly analyses that go beyond these topics.⁴⁴⁸ In their analyses, many scholars have noted the themes of *moral* fallibility, which are ubiquitous not only in *No. 10* but also in much of the rest of Madison's constitutional thought, as well as that of the rest of the American founding fathers.⁴⁴⁹ However, I have found no scholarly work that notes the central role played by what I have called chronic *epistemic* fallibility in Madisonian thought. Yet, I believe the notion of chronic epistemic fallibility plays a crucial role in Madison's arguments in *No. 10* and potentially bears implications for many of the arguments he posed in other *Federalist Papers*. Madison's view of humankind as chronically *epistemically* fallible is integral to his defense in *No. 10* of a large republic both in terms of *geography* and *population*. It also plausibly informs his support for the "checks and balances" he discusses in later *Federalist Papers*. In both cases, I will argue, Madison's thoughts imply that the proper way for political theorists and/or practitioners to respond to the problem of chronic epistemic fallibility is to adopt policies that foster *perspectival diversity*.

To make this case, let me first illustrate the ways Madison demonstrates an appreciation for epistemic fallibility in *Federalist No. 10*. The most obvious example comes when Madison is arguing why, despite its obvious dangers, it is inadvisable to attempt to get rid of faction entirely. Madison argues that doing so could be accomplished only by one of two methods, "the one, by

446E.g. *Storer v. Brown* (1974), *California Democratic Party v. Bill Jones* (2000)

447*Nixon v. Shrink* (2000)

448Political theorist Robert Dahl, for instance, argues against Madison's arguments in both *No. 10* and *No. 51* on the grounds that his claims about the dangers of majority rule are empirically unsupported and that the effectiveness of the "checks and balances" called for by Madison ultimately stems from their being enforced through public accountability, rather than being a separate check on the power of politicians (Dahl 1956, 1-33). Elsewhere, Laurence E. Lynn, Jr. argues that the gridlock commonly complained of by those frustrated with American politics was foreseen and viewed by Madison as an acceptable side-effect of the separation of powers (Lynn 2011).

449While the most famous examples of Madison's emphasis on human beings' moral fallibility are found in *No. 51*, wherein Madison penned the now-immortal phrase, "If men were angels, no government would be necessary," they can also be found in his other *Federalist Papers*, such as when he insists in *No. 10* that "[s]o strong is [the] propensity of mankind to fall into mutual animosities, that where no substantial occasion presents itself, the most frivolous and fanciful distinctions have been sufficient to kindle their unfriendly passions and excite their most violent conflicts," and indeed "neither moral nor religious motives can be relied on" to keep a determined majority faction from depriving its rivals of their rights."

destroying the liberty which is essential to [faction's] existence; the other, by giving to every citizen the same opinions, the same passions, and the same interests.”⁴⁵⁰ Madison rejects the first course of action because the “cure” of depriving people of this sort of liberty is “worse than the disease” it is meant to treat. Thereupon follows one of the most famous quotes in *No. 10*: “Liberty is to faction what air is to fire.” As famous as this passage is, though, Madison's explanation *why* liberty so inevitably leads to faction has received far less attention. *Why*, specifically, should we expect men under conditions of liberty to divide into factions? *Why* does Madison say that it is just as foolish to try give everyone the same *opinions* as to try to give them the same “passions and interests?” The answer to both questions can be found in the next paragraph, where the very first reason Madison gives for why it is “impracticable” to give everyone the same opinions, passions and interests is that “[a]s long as the reason of man continues to be *fallible*, and he is at liberty to exercise it, different opinions will be formed.” In the next sentence Madison notes that because of the tight connection between man's “reason and his self-love,” each person's station in life is bound to influence his point-of-view. Because of the “diversity in the faculties of men,” writes Madison, “the possession of different degrees and kinds of property immediately results; and from the influence of these on the sentiments and views of the respective proprietors, ensues a division of the society into different interests and parties.” In other words, men's different societal positions are likely to give them a different, and imperfect, *perspective*, leading them to make different, and imperfect, *judgments* about what is true. Because these imperfections will lead their judgments to diverge from one another, we should expect men left free to come to conclusions on the basis of their own judgment to disagree with one another, and accordingly to divide into different interest groups to promote different causes and champion different public policies.

Though Madison himself did not emphasize the connection and it is therefore impossible to say how significant a role it played in his thinking, this notion of man as chronically *epistemically* fallible can serve as the starting point for powerful *epistemic* defenses of some of Madison's core political commitments, particularly his defense of representative government and federalism. Epistemic support for the “checks and balances”⁴⁵¹ of *No. 51* on can be inferred from

⁴⁵⁰*Federalist No. 10*, 53, my emphasis

⁴⁵¹I put the term “checks and balances” in quotes because, while the phrase has become synonymous with and encapsulates a key feature of Madisonianism, the term itself was not used by Madison.

his comment in *No. 10* that “No man is allowed to be a judge in his own cause, because his interest would certainly bias his judgment.”⁴⁵² On this view, institutions like judicial review, the Presidential veto, and the division of the legislature into separate bodies elected by distinct populations may be defended not *just* because “neither moral nor religious motives can be relied on”⁴⁵³ as adequate means of preventing any unified body from becoming “the very definition of tyranny” if allowed to accumulate “all powers, legislative, executive, and judiciary” in its hands,⁴⁵⁴ but also because the *judgment* of any single group of decision-makers is likely to be *skewed* in favor of its own interests. The implication would be that the function of Madison's signature “checks and balances” is not *just* to encourage each branch of government watch out for *corruption* by making “ambition counteract ambition,” but also to make sure that the *decisions* reached by any single group of decision-makers, who are animated by a given set of interests and therefore whose judgments are likely to be skewed in a given direction, are subjected to *external checks* by a group whose interests are different from the first and therefore whose judgments, while equally fallible, are likely to be skewed in a *different direction*. Another way of phrasing this might be to say that this reading of Madisonianism sees the benefits of “checks and balances” as stemming not just from their tendency to protect the public from becoming prone to the vagaries of a single individual group's *will* by counteracting it with *other opposed wills*, but also from their tendency to protect the public from becoming prone to the vagaries of a single individual or group's *judgment* by forcing the decision-making process to pass through the judgment of multiple individuals/groups who have a *different perspective*.

As I said above, it is not clear whether this notion of chronic epistemic fallibility played a significant role in endorsing “checks and balances” for Madison himself, though I would argue that both the general prevalence of epistemic themes in *No. 10*⁴⁵⁵ and the priority epistemic concerns are given whenever they are included in a list of multiple reasons provided in support

⁴⁵²*Federalist No. 10*, 53

⁴⁵³*Ibid.*, 54

⁴⁵⁴*Federalist No. 47*

⁴⁵⁵In addition to being central to Madison's explanation for the very existence of faction and his opposition to allowing anyone to be judge in their own cause, Madison also argues in favor of federalism in *No. 10* by insisting that each representative being “chosen by a greater number of citizens” will lead to “men who possess the most attractive merit and the most diffusive and established characters” being “in the large rather than the small republic” and in favor of representative rather than popular democracy partly on the grounds that it substitutes representatives with “enlightened views” that help to “render them superior to local prejudices.”

of a given argument for in *No. 10* suggests it was at the forefront of his mind when writing the latter essay at least.⁴⁵⁶ My point in bringing Madison into the conversation is not primarily to say that his defense of the “checks and balances” he discusses in *Federalists No. 47-52* is based largely on his chronically fallible view of human epistemology. Instead, my point has been to show that strong epistemic arguments in favor of such “checks and balances” are implicit in Madison's chronically fallible conception of human epistemology, whether he made the connection or not. What I now want to do is (a) show that Madison's response to what he viewed as humankind's chronic *moral* fallibility was to emphasize the importance of dividing decision-making authority between groups holding a *diversity of interests* and (b) argue that my epistemic reading of Madison suggests the proper response to humankind's *epistemic* fallibility is to inscribe into the Constitution measures that foster *perspectival diversity* among the public.

Task (a) is straightforward. Madison's *Federalist Papers* are peppered with arguments that respond to chronic moral fallibility by recommending policies designed to foster a diversity of interests within the relevant decision-making body. It is, of course, the essence of his argument that “ambition must be made to counteract ambition” by incorporating “checks and balances” into the constitution of a government, which amounts to relying on the *diversity of interests* between different branches of government to serve as a “check” on the moral fallibility of any single branch. But Madison's defense of federalism also relies on this kind of reasoning. He opposes “pure democracy” because he takes it for granted that a “*common* passion or interest will, in almost every case, be felt by a majority of the whole,” inducing them to pursue their own interest even if that means sacrificing the “weaker party or an obnoxious individual.” This worry forms the basis of one of his arguments in favor of the republican form of government, since “the greater number of citizens and extent of territory which may be brought within the compass of republican than of democratic government” allows it to “take in a greater variety of parties and

⁴⁵⁶Not only does Madison stress chronic epistemic fallibility in numerous ways through the essay, but each time he lists multiple considerations in support of a given claim, epistemic reasons appear first. As mentioned above, it is his first explanation for why a diversity of opinions is inevitable under conditions of liberty. It is also the first reason he gives in *No. 10* for opposing any situation where a man should be “allowed to be a judge in his own cause.” Only afterward does he add that such a condition might also “not improbably, corrupt [the judge's] integrity.” Later, the first reason he gives for insisting it is “vain to say enlightened statesmen” will be able to mitigate the threat of faction is that “[e]nlightened statesmen will not always be at the helm,” only afterward going on to mention the improbability of “indirect and remote considerations” prevailing over “the immediate interest which one party may find in disregarding the rights of another or the good of the whole.”

interests,” making it “less probable that a majority of the whole will have a common motive to invade the rights of other citizens; or if such a common motive exists, it will be more difficult for all who feel it to discover their own strength, and to act in unison with each other.” Thus, just as “checks and balances” serve to combat politicians' moral fallibility by introducing a diversity of interests into the decision-making process at the level of government, Madison envisions the sort of diversity of interests and passions likely to be encompassed by a larger, more geographically widespread republic to be a safeguard for minorities against the moral fallibility of potentially tyrannous majorities. The argument for point (b) is similar, only this time the type of fallibility dealt with is *epistemic* and the type of diversity seen as the proper safeguard is *perspectival*. On the epistemic reading of Madisonianism articulated above, Madison's defense of “checks and balances” would be read as safeguarding against the epistemic fallibility of one branch of government by making their decisions answerable to a different branch with a different perspective, while Madison's defense of a larger republic would be read as safeguarding against the epistemic fallibility of any single group of citizens by requiring political leaders and policy initiatives to appeal to a wider variety of different groups, all of whom view things from a different perspective. This reading suggests Madison's recognition of humanity's chronic *epistemic* fallibility led him to emphasize the need to inscribe *perspectival* diversity into political decision-making processes and to advocate a larger and more geographically widespread republic as a means of incorporating a greater degree of perspectival diversity among the voting public. Madison treated the fostering of diversity as the appropriate response to *both* moral fallibility *and* epistemic fallibility. If we take Madison to have envisioned humankind as being chronically fallible in both senses, as I believe he did, we can see how central the desire to foster a diversity of interests and perspectives in the formation of our political decision-making processes was to Madisonian thought.

Now, it is true that, like the other American founding fathers, Madison at the time of *No. 10*'s writing was an elitist and an anti-democrat with a low opinion of the cognitive faculties and likely motivations of the poor.⁴⁵⁷ These sentiments may well be behind Madison's comment in

⁴⁵⁷These opinions may, however, have changed after the writing of *No. 10*, due in large part to Madison's observations of the early American republic. So, at least, argues Dahl, who in an essay published nearly 50 years after his initial critique of Madisonianism (cited above) writes that his research on Madison in the intervening years led him to believe that “experience with the rapidly emerging American democracy led James Madison to views that I would regard as somewhat more democratic than those he expressed at the Constitutional

No. 10 that the “diversity in the faculties of men” is the main reason why some acquire property while others do not. Nevertheless, the text of *No. 10* itself makes no explicit endorsement of any particular point-of-view as being *more* immune from this general condition of fallibility than any other. It is the reason of *man* that is fallible and *man's* self-love that causes his opinions and passions to wield a “reciprocal influence on each other,” causing each person's judgment to be skewed. Though Madison may himself have been an elitist, his words in *No. 10* portray epistemic fallibility as a *universal* human trait, one that is bound to skew the reasoning of *every* individual human being in a way that reflects their own interests, upbringing and station in society. It is striking that Madison who, like most of his property-owning contemporaries, was largely unrepentant in his elitism, never comes close to suggesting that any single group, whether it be propertied elites, elected representatives, or any other group is immune, or even appreciably resistant, to either moral or epistemic fallibility. Instead, he writes throughout as if *all* groups' judgment is imperfect and skewed by its own interests and *all* groups are susceptible to corruption. This is likely why Madison's response to the problems of fallibility was to emphasize bringing a *diversity* of interests and perspectives to bear in public decision-making processes, rather than trying to identify a *particular* group whose interests were presumed to be benevolent and whose perspective is presumed to be the “right” one. Thus, ironically, though Madison was unabashedly elitist and explicitly opposed “pure democracy,” his thinking never comes close to endorsing anything like the level of routine deference to established authorities characteristic of the writing of Progressives, most of whom view themselves as ardent defenders of democracy and opponents of the sort of elitism embraced by Madison. This difference stems from the difference between Madison's and Progressives' treatment of human fallibility. Madison treated both moral and epistemic fallibility a *chronic* conditions that apply to *all* human beings and institutions while Progressives treat them as conditions that apply primarily to *lay citizens* but do not merit much concern when it comes to their favored institutions.

I believe Madison's approach is superior, but before I go my reasons let me first explore the thinking of another founding father. As it turns out, Madison was not the only member of his generation whose notion of fallibility as a chronic condition led him to emphasize the need to foster perspectival diversity and abjure anything resembling the sort of deferential model of

Convention” (Dahl 2005, 439).

public policy-making and civic epistemology emphasized by Progressives. Another was Thomas Jefferson, whose thoughts on the proper meaning and role of “the press” and “freedom of the press” I will now explore.

Diversity as the Proper Response to Chronic Fallibility in Jefferson's Thoughts on “The Press”

Throughout this dissertation I have used the legacy news media as a running example of Progressives' approach to civic epistemology. I did not choose that example at random. I believe Progressives' treatment of the legacy news media, its role in democracy, and the attitudes citizens ought to take toward it, as perhaps the single best encapsulation of what is wrong with the way Progressives, and I believe most citizens, approach civic epistemology. I have criticized that approach as excessively deferential and uncritical, but I have not yet subjected Progressives' views of the meaning and proper role of “the press” and “freedom of the press” to focused criticism. I will do so now, arguing that the notions of “the press” and “freedom of the press” have in contemporary times become warped in a way that encourages us to overrate the importance of the legacy news media and underrate the benefits of 21st century communications technologies to the epistemic well-being of democracy. A better way, I will argue, is exemplified in the writings of Thomas Jefferson, which are grounded in the same conception of chronic fallibility as I argued in the previous sub-section to be central to the writing of James Madison.

To see the value of Jefferson's approach, and how it differs from that of Progressives, it is first necessary to cobble together the single coherent view of “the press,” its democratic value, and its role in society that lies behind the Progressive positions I have emphasized at various points in this dissertation, a view which I believe to be exceedingly common among contemporary democratic citizens in America and many other parts of the world and to be implicit in much of our secondary and post-secondary civics curricula. For Progressives, as well as for most contemporary Americans and I suspect many democratic citizens around the world, the term “the press” is a synonym for “the media,” which is in turn a near-synonym for “the news media,”⁴⁵⁸ or what I have dubbed the “legacy news media.” That is to say, the term “the press” means a professionalized group of information-publishers who work for established institutions like *The New York Times*, *CBS*, and *The Economist*. The American First Amendment

458Ladd (2012)

famously guarantees “freedom of the press,” and such a guarantee is viewed by many to be an indispensable aspect of any democracy worth defending.⁴⁵⁹ Given prevailing ideas of what “the press” is, it should not be surprising that “freedom of the press” has come for many to mean “freedom of the legacy news media.” Accordingly, when contemporary political commentators speak about the democratic importance of “freedom of the press,” they nearly always speak as though the sole or primary democratic function of that freedom is to ensure that the legacy news media can perform what are usually taken to be its main democratic tasks. These are, in the words of several of the Progressives cited previously in this dissertation, to “serve as a credible watchdog,”⁴⁶⁰ “ferret out crucial political truths,”⁴⁶¹ “keep ‘false facts’ and demagoguery at bay,”⁴⁶² and serve “as an alternative, and antidote, to the contamination of the information environment” by misinformation.⁴⁶³ This view of the primary function of “freedom of the press” as intended to allow professional journalists to carry out a unique and specialized democratic task is also implied when the non-profit organization Freedom House writes that a “free press can inform citizens of their leaders’ successes or failures” and “convey the people’s needs and desires to government bodies,” both of which phrases clearly imply that “the press” is some institution *other* than “the people” and that the importance of “freedom of the press” is related to enabling that institution to perform its proper role. During the Trump Presidency, then-President Trump was often accused of endangering “freedom of the press” by undermining people’s trust in the legacy news media,⁴⁶⁴ banning certain members of the legacy news media from press conferences,⁴⁶⁵ and insinuating on his Twitter feed that he would look into revoking the licenses of several networks whose criticism of him he viewed as unfair.⁴⁶⁶

The notions of “the press” and the importance to democracy of maintaining “freedom of the press” implicit in these examples are as obvious as they are familiar. “The press” means a

459The reader may suppose I myself endorsed such a view in Chapter One, but what I actually defended there were what I called “core liberal freedoms of communication, among which are freedom of speech and freedom of the press.” As I will shortly argue, my notion of “freedom of the press” and its value to democracy diverges so starkly from what is implied by the most common contemporary interpretations of that expression that I view it as an entirely different thing, built on an entirely different conception of civic epistemology.

460Jamieson (2018), 218

461Goldman (2009), 3

462Schiffrin (2017), 123

463UNESCO (2018), 9

464*Boston Globe* (2018)

465Pilkington (2018)

466Rafferty (2017)

professionalized group of journalists, i.e. the legacy news media, and freedom of the press primarily means enabling that professionalized group to do what the legacy news media has traditionally done, such as print “newsworthy” stories, publish the words and deeds of power-holders and attend and report on press conferences. Also apparent in several of the examples given is the view held by many commentators that public trust in the legacy news media is important to the well-being of democracy and that public distrust is a bad sign, as I pointed out in Chapter Two.

As familiar as these views are, they are surely an extraordinary departure from the notions of “the press” and the role of “freedom of the press” the American founding fathers had in mind when they included press freedom among the rights guaranteed by the First Amendment. The most obvious reason this must be the case is that there was no legacy news media, nor anything remotely resembling it, at the time the First Amendment was composed. The contents of the era's papers resembled the contemporary *New York Times* and *Wall Street Journal* only in the sense that they were written communications involving headlines and body text. None of the other conventions that typify contemporary legacy journalism were present. There was no separation of “news” from “opinion,” nor was any attempt made to clearly separate advertisements from factual reports of perceived public importance. As Schudson reports, newspaper publishers in the colonial period made no affirmative effort at “newsgathering” whatsoever, a process that would not be viewed as part of a newspaper's operations “for about a century to come.”⁴⁶⁷ Publishers' notions of the sort of “news reporting” they took to be their proper province was summed up in the words of Benjamin Harris, editor of America's first newspaper, as consisting of passing on “such considerable things as have arrived unto our Notice,” which Barber anticipated requiring the publication of a single paper about once a month, though he conceded that more publications might be required “if any Glut of Occurrences happen.”⁴⁶⁸ The composition of colonial-era newspapers was typified by the publications of Benjamin Franklin and his brother James, which were a hodge-podge of rumors from recently arrived sailors, advertisements from local merchants, and “letters” from concerned readers, many of which were made up by the more famous Franklin and published under

467Schudson (1990), 19

468Boorstin (1960), 7

pseudonyms like “Silence Dogood” and falsely presented to readers as stemming from members of the community. A good example is the November 22nd edition of the *New York Daily Advertiser*, where an immortal piece of political theory, Madison's *Federalist No. 10*, is joined in a single page by the reprinted political tirade of an unnamed editoralist at an unspecified “Liverpool paper,” a smattering of rumors passed on by passengers on the recently-arrived ship, the *Hector*, and the melodramatically relayed account of a condemned prisoner who had (supposedly) been found floating in the local bay by a visiting sailor on his way to being exported to “Old France, and there shot to death by way of example” for “some misbehaviour” when, “impelled by that natural love of life which is implanted in our nature, he took the resolution of getting out of one of the port-holes, and dropping into the water, chosing [*sic*] to run the risque of drowning, rather than proceed to his native country, where certain death awaited him.”⁴⁶⁹

As Jonathan Ladd writes, because this is what “the press” was at the time of the writing of the American Constitution, when the American founding fathers defended “freedom of the press,” they could not have had “in mind protections for a profession called journalism, because such a profession did not exist.”⁴⁷⁰ To them, “the press” did not refer to a specialized profession dedicated to a democratically important mode of information-gathering and dissemination. Instead, it meant a jumbled smorgasbord of advertisements, fiction, editorials and rumors passed on from thinly vetted sources. Another way of saying this would be to say that “the press” the American founding fathers saw themselves as defending looked, in terms of accuracy, content and organization, a lot more like a Facebook or Twitter feed, or a click-bait filled news feed on a “new media” website, or even an edition of *National Enquirer*, than it did an edition of the contemporary *New York Times*.

Progressives would surely be appalled if they woke up tomorrow to find *The Washington Post* peddling the contents regularly published by the newspapers of the colonial era. So, indeed, were many of the founding fathers, not least Thomas Jefferson, whose collected letters make it clear that he considered the average contents of the press of his time to be of exceedingly poor quality. “As for what is not true,” he writes in one letter, “you will always find abundance in the

⁴⁶⁹*New York Daily Advertiser* (Nov. 22nd, 1787)

⁴⁷⁰Ladd (2012), 21

newspapers.”⁴⁷¹ In another he opines that “Nothing can now be believed which is seen in a newspaper. Truth itself becomes suspicious by being put into that polluted vehicle.”⁴⁷² These quotes illustrate that Jefferson viewed the press of his day to frequently be an unreliable source of information. But its vices were not limited to inaccuracy. In another missive, he blames “the press” for thriving on sensationalism and for sowing division in society, accusing the press of living “by the zeal they can kindle & the schisms they create.”⁴⁷³ In yet another he claims that “[t]o divide those by lying tales whom truths cannot divide, is the hackneyed policy of the gossips of every society.”⁴⁷⁴

These complaints about the factual unreliability and divisiveness of the colonial-era press bear a striking resemblance to contemporary Progressives' complaints about the 21st century informational environment, which, as we have seen, focus heavily on its role in facilitating the spread of misinformation and exacerbating the growth of polarization. More striking, however, are the differences. First, Jefferson viewed “the press” as a major *contributor* to these problems, whereas Progressives primarily portray these as problems that have been exacerbated by the rise of *alternatives* to “the press,” and view the legacy news media as a bastion of accurate, largely balanced, information and societal cohesion.⁴⁷⁵ One might be tempted to write this difference off as mere semantics given what I have said so far about how the founding fathers' conception of “the press” differed from Progressives were it not for a second difference. This is that, while Progressives' response to the threats of misinformation and polarization fomented online has been to portray them as *threats* to democracy and accordingly to try to impose some sets of authority to *control* them, Jefferson viewed “the press” he so roundly condemned on these very same grounds as an *indispensable asset* to democracy and staunchly *resisted* any suggestion that they ought to be controlled by any over-arching authority. “[O]ur liberty,” he wrote in 1786, “cannot be guarded but by the freedom of the press, nor that be limited without danger of losing

471 Letter to Barnabas Bidwell, July 5th, 1806. All Jefferson correspondence cited in this dissertation are freely available in digital form and may be found by searching the recipient's name and date in the United States National Archives at Founders Online, <https://founders.archives.gov>.

472 Letter to John Norvell, June 11th, 1807

473 Letter to Elbridge Gerry, March 29th, 1801

474 Letter to George Clinton, December 31st, 1803

475 The latter point is implicit in Sunstein's view of legacy news organizations as capable of providing “shared experiences for a heterogeneous public,” which he views as something of an antidote to the danger of polarization (Sunstein 2017, 43).

it.”⁴⁷⁶ Two decades later he called freedom of the same press he viewed with such scorn and derision the “most effectual” means yet discovered of ensuring that man is “governed by reason and truth.”⁴⁷⁷ Two decades later still Jefferson wrote to Lafayette that freedom of the press whose accuracy he disparaged and whose divisiveness he condemned was the “only security of all.”⁴⁷⁸

The differences between the Jeffersonian view of “the press” and that of contemporary Progressives, then, cannot be reduced entirely to the semantic gap that arose between the late 18th and early 21st centuries regarding the meaning of the term “the press.” Instead, they indicate a fundamental difference in the epistemic outlook of Jefferson vis-a-vis that of Progressives. Jefferson's defense of freedom of “the press,” despite holding such apparent scorn for its usual performance, was based on the same fallibility-centric worldview as Madison's defense of federalism and “checks and balances.” The numerous quotes given above clearly demonstrate that Jefferson was as well aware of the fallibility of “the press” of his era as are Progressives of the fallibility of the social media users of our own. Yet, Jefferson refused to respond to that fallibility by recommending the imposition of censorship by “fact-checkers” or any other authority because his recognition of publishers' fallibility was balanced by an equally clear view of the fallibility of any authority that might take up such a censorial role. Again, this sense of fallibility included both moral and epistemic components. Sounding very Madisonian, he assumed that any government given the power to exercise censorship would come to employ a “standing army of newswriters who, without any regard to truth or to what should be like truth...put into the papers whatever might serve the minister.”⁴⁷⁹ This staunch defense of freedom of the press as a means of staving off corruption also drew support from epistemic considerations. He took it for granted that epistemic fallibility makes it “so difficult to draw a clear line of separation between the abuse and the wholesome use of the press” and to discriminate “between truth & falsehood” as to undermine the very idea that any authority could ever be fit to do the job.⁴⁸⁰ Like Madison, Jefferson saw both moral and epistemic fallibility everywhere. They were not the unique or special condition of any particular group of people, nor was there any form of training that could rid any individual or group of it to the degree that

476Letter to John Jay, January 25th, 1786

477Letter to John Tyler, June 28th, 1804

478Letter to Lafayette, November 4th, 1823

479Letter to G. K. van Hogendorp, October 13th, 1785

480Letter to M. Pictet, February 5th, 1803

would justify their being treated as an authority to which routine deference could safely be granted. Instead, they were assumed to be a fundamental, and chronic, aspects of the human condition, common to lay citizens, newspaper editors, government officials, and everyone else, and therefore bound to infect all institutions created by and composed of such human beings.

Just as Jefferson's view of the human condition mirrored Madison's in being infused by the assumption of chronic fallibility, his proposed method of addressing the problems caused by fallibility was the same. Since no particular authority could ever be trusted to be either epistemically infallible or incorruptible, the excesses of "the press" could not be solved by imposing some set of authorities to censor its contents or by encouraging lay citizens to adopt a disposition of skepticism toward "the press" while placing its trust in the claims of some other information provider it could safely presume to be immune, or largely immune, from such flaws. Instead, like Madison, Jefferson responded to chronic fallibility by trying to ensure that the informational environment was set up in such a way as to *facilitate diversity*. In Jefferson's words, the best means of protecting "truth and reason" was *not* to identify some single source, or narrow range of sources, who were their guardians, but instead to make sure the public had regular access to "a full hearing of all parties."⁴⁸¹ When Jefferson wrote of his aspiration that "man may [come to] be governed by reason and truth," his proposed means of pursuing this ideal was not to assign the task of figuring out what "reason and truth" demanded to any specialized community but instead to "*leave open* to [the public] *all avenues* to the truth."⁴⁸² Thus, despite his scornful attitude toward "the press," Jefferson's solution to its fallibility was neither to do away with it nor to encourage the public to view the press with skepticism while placing its trust in some other set of authorities. In this way, Jefferson's approach to "the press" and the proper role of "freedom of the press" closely resembled Madison's approach to the proper formation and scope of the government and the republic. In each case, chronic fallibility was assumed to be a deep, serious and intractable condition whose proper solution was not to impose and/or advocate widespread deference to the "right" authorities but rather to try to foster conditions that allowed for, or even actively encouraged, the emergence of *a diversity of interests and perspectives* and their injection into political decision-making processes and public discourse.

⁴⁸¹Jefferson (1805)

⁴⁸²Jefferson to John Tyler, June 28th, 1804, my emphasis

I call the approach taken by Madison and Jefferson, which treats fallibility as a *chronic, universal* condition properly addressed by seeking to foster diversity, *prophylactic pluralism*. While pluralism is a familiar concept for liberals, much of the discussion entailing that concept has revolved around whether and to what extent pluralism is a *good in itself*, or at least something that ought to be “tolerated.”⁴⁸³ Historically, the discussion has typically focused on arguments about just how wide a range of *values* and *lifestyles* can and ought to be integrated into society. More recently, however, some political theorists have invoked the notion of pluralism in a different way, defending pluralism not only as a good in itself but also on the grounds that it is likely to produce beneficial *outcomes*, both politically and epistemically.⁴⁸⁴ I believe the reading of Madison and Jefferson I have offered here supports the latter line of argument, defending pluralism on the grounds of its practical consequences for collective knowledge and oversight of authorities.

Epistemic Pluralism and Modest Optimism about Lay Citizens' Judgment in Madison and Jefferson

What I have described as the prophylactic pluralism of Madison and Jefferson had two elements, which can be thought of as *moral* and *epistemic*. The moral element saw people as chronically *morally* fallible and sought to combat this flaw by incorporating a diversity of *interests* into the relevant decision-making bodies. This aspect of the thinking of the American founding fathers has been well documented and extensively discussed. The epistemic element, which saw people as *epistemically* fallible and sought to combat this flaw by incorporating a diversity of *perspectives* into the relevant decision-making bodies, has not. This approach I call *epistemic pluralism*, and will be the focus of the rest of my commentary here.

Neither Madison nor Jefferson explicitly articulated the view I have described as epistemic pluralism, but I believe this was simply because it was the outlook of many of the founding fathers, as well as many of the citizens that inhabited the United States in the 18th and

⁴⁸³Locke's *Letter on Toleration* is of course the foundational text here. One prominent version of the long-standing debate about the virtues and limits of this sort of pluralism has played out between “political liberals,” such as John Rawls, and communitarians like Charles Taylor and Michael Sandel. For an excellent argument about why this focus on toleration is excessively narrow and constricting, see Muldoon (2016).

⁴⁸⁴For a series of arguments focused on the *epistemic* benefits of diversity, see Page (2007). For both epistemic and general political benefits, see Landemore (2012).

19th centuries, and therefore was simply assumed as common ground that did not need to be explicitly defended. This is most evident in three features these thinkers' political thought had in common. First, chronic fallibility was always *at the forefront* of both thinkers' political writing. Second, their response to this chronic fallibility was not to advocate deference to any single source, institution, research methodology, or any other form of authority presumed to have a superior grasp on the truth, but instead to foster conditions likely to encourage the bringing to bear of *multiple, differently-situated perspectives* on public problems. Third, their view of human judgment, notably that of lay citizens, was infused with the same sort of *modest optimism* implied by epistemic pluralism, and informed their emphasis on the value of fostering different perspectives. I have already dwelt at length on the attention paid to chronic fallibility by Madison in his *Federalist Papers* and Jefferson in his letters about the press and its proper societal role, as well as their shared preference for dealing with that fallibility by emphasizing the need to make sure a diversity of perspectives and interests had an opportunity to influence public decisions. What remains to be demonstrated is the way these recommendations were also supported by both thinkers' *modestly optimistic* view of the judgment of lay citizens. To that task I now turn, beginning with Madison.

As the passages quoted in Section Two clearly demonstrate, Madison was no romantic, either about human nature nor about the lay citizenry of the United States. He opposed “pure democracy” because he assumed that, no matter where or in what era it occurred, as soon as a “common passion or interest” came to be “felt by a majority of the whole,” no force could be expected to stop them from oppressing “the weaker party or an obnoxious individual.” Nothing in his *Federalist Papers* suggests he held any more romantic a view of the public's wisdom than he did of its capacity for self-restraint. On the contrary, as argued above, in *No. 10* he assumes all men's reason to be both fallible in general and to be skewed by their “self-love.”⁴⁸⁵ Meanwhile, his discussion in *Nos. 47-52* revolves almost entirely around ways of getting different members of government to watch over *each other*. In *No. 51* he gestures briefly toward the public when he writes that a “dependence on the people is, no doubt, the primary control on the government,” but he follows even that brief concession with a the clarification that “experience has taught

⁴⁸⁵*Federalist No. 10*

mankind the necessity of auxiliary precautions.”⁴⁸⁶ Madison's opposition to “pure democracy” and endorsement of a trustee rather than a delegate model of representation in *No. 10*, as well as his recognition of the need for “auxiliary precautions,” dwelt on at length in *Nos. 47-52*, clearly indicate that his faith in the ability of lay citizens to support sound public policies and keep watch over public officials was limited.

Part of the reason Madison's views of the lay public were far from romantic almost certainly came from the elitism he shared in common with many other founding fathers as members of the landed gentry. Part of it, too, seems to have come from his assumption that the quality of individuals' judgment varied from person to person, partly due to their innate characteristics,⁴⁸⁷ partly due to training,⁴⁸⁸ and partly due to the advantages and disadvantages of different individuals being placed in different societal stations.⁴⁸⁹ Since many citizens would neither be especially bright, attain a high level of intellectual training, or hold a position that put them in a good position to form an accurate view of the general state of the republic, Madison clearly did not view lay citizens as capable of being anything remotely close to “omniscient.”⁴⁹⁰ Yet, despite his elitism and despite his view of citizens as beset by relatively severe epistemic limitations, Madison nevertheless saw an integral role for the judgment of lay citizens in the republic he envisioned, a role which implied a certain modest optimism about the judgment of just those lay citizens. Though he may have written a great deal about the importance of “auxiliary precautions” in keeping the various branches of government in line, Madison still insisted that there was “no doubt” that the *primary* “control on government” was its “dependence on the people,” an insistence that informed his clear preference for the legislature to be the strongest branch of government.⁴⁹¹ Madison also illustrates a modest optimism about lay judgment and implies a potential affinity for epistemic pluralism when he endorses federalism on the grounds that in a large republic “each

⁴⁸⁶*Federalist No. 51*

⁴⁸⁷“The diversity in the faculties of men” are sufficiently great that they are “not less an insuperable obstacle to a uniformity of interests” than their general fallibility, writes Madison in *No. 10*.

⁴⁸⁸I take this to be the essence of Madison's claim that “peculiar qualifications” are an “essential” feature of the judiciary in *No. 51*.

⁴⁸⁹Elected representatives, for example, are viewed by Madison as in a position to “refine and enlarge the public views” and “best discern the true interest of their country” in *No. 10*.

⁴⁹⁰The term “omniscient” here is a reference to Lippmann, who criticized democratic thinkers of his era of holding a naively romantic view of citizens (Lippmann 1925, 21).

⁴⁹¹“In republican government, the legislative authority necessarily dominates” (*Federalist No. 51*).

representative will be chosen by a greater number of citizens,” meaning that “it will be more difficult for unworthy candidates to practice with success the vicious arts by which elections are too often carried; and the suffrages of the people being more free, will be more likely to centre in men who possess the most attractive merit and the most diffusive and established characters.”⁴⁹² Madison's view of *the people* as the *primary* control on government and his belief that a large, diverse republic composed of citizens belonging to “a greater variety of parties and interests” is likely to result in *both* the selection of better representatives *and* resistance to abuse of minorities by a single overwhelming power both reflect a modest optimism about lay citizens' judgment. Though many *individuals'* judgment is bound to be deficient, *on average* their judgment is *good enough* that, as a collective, they may be trusted to serve as the primary line of defense against tyranny and to select good (enough) leaders. The larger the number of citizens and the wider the variety of “parties and interests” – in other words, the greater number and breadth of *perspectives* – allowed to participate in these processes, the better.⁴⁹³ Thus, it was not just his view of humankind as *chronically fallible* that drove Madison to emphasize the importance of facilitating diversity but instead the *simultaneous* endorsement of that view with a certain *modest optimism* about the judgment of that same chronically fallible species. His subscription to this combination of views is why I believe Madison's epistemic outlook was something very like the sort of epistemic pluralism I described in Section Three.

So too, I believe, was Jefferson's. It is true that, compared to those of Madison, Jefferson's depictions of lay citizens were often famously romantic. Romantic though they may have been, however, they were informed by an awareness of the public's fallibility that tempered these romantic excesses. Jefferson admitted that the public regularly made “errors,” both in terms of their ability to tell fact from fiction on a case-by-case basis⁴⁹⁴ and in terms of their evaluations

⁴⁹²*Federalist No. 10*

⁴⁹³It goes without saying that Madison was a slave owner who, like most of his generation, only wished for property owning white males to participate in elections. While these facets of Madison's record cannot be ignored and surely mar his legacy while truncating the breadth and number of perspectives allowed to take part in the elections of the early republic, his arguments in *The Federalist Papers* are written in terms that make them readily compatible with our more cosmopolitan and inclusive views at present. Indeed, I believe the co-incidence of the improvement of our laws, on average and over time, while the electorate has expanded and become more inclusive vindicates Madison's writings in *The Federalist Papers* even as it rebuts the public policies he supported in practice.

⁴⁹⁴Letter to G. K. van Hogendorp, Oct. 13th, 1785

of political candidates.⁴⁹⁵ Jefferson viewed the public as fallible in the same way he viewed newspaper editors and political leaders as fallible. Indeed, the letter to Stuart shows that he took it for granted not only that individual citizens would often adopt mistaken beliefs but that on occasion, perhaps not even infrequently, the entire *public* could be deceived. In this way, Jefferson would have agreed with contemporary Progressives. Still, recognition of the limitations and error-proneness of lay citizens notwithstanding, Jefferson would nevertheless have adamantly disagreed with the Progressive tendency to respond to this possibility by looking to impose some sort of authority on the informational environment, either through censorship by “fact-checkers” or by advocating for citizens to adopt an attitude of routine deference to institutions it could safely assume were guardians working “in the service of truth”⁴⁹⁶ helping to “keep 'false facts' and demagoguery at bay.”⁴⁹⁷ Instead, as we have seen, rather than proposing the imposition of, or routine deference to, authority, Jefferson advocated responding to the inevitable danger of much of the public being deceived either deliberately or accidentally by the things they read in “the press” by jealously guarding that very same press' *freedom*.

How could Jefferson, who viewed “the press” at the time in much the same way Progressives today view social media and the blogosphere – that is, as so untrustworthy that “[N]othing can now be believed which is seen in a newspaper,” making truth itself become “suspicious by being put into that polluted vehicle”⁴⁹⁸ – also depict it as the “only tocsin of a nation?”⁴⁹⁹ The key can be found in his *modest optimism* about the lay public's judgment. Even though Jefferson admitted that the lay public could, and often would, be deceived by the information passed on in those papers, *on average and over time* he believed their judgment was reliable enough to allow them to correct their errors. Thus, while the people “may be led astray for a moment,” they “will soon correct themselves.”⁵⁰⁰ Thus, in stark contrast to Progressives, the staunchness of Jefferson's defense of “freedom of the press” did *not* stem from Jefferson's view of “the press” as an institution whose established procedures and training imparted its contents with a high degree of reliability and its practitioners a high degree of trustworthiness. Indeed, his

495 Letter to Edward Carrington, January 16th, 1787

496 UNESCO (2018), 9

497 Schiffirin (2017), 123

498 Letter to John Norvell, June 11th, 1807

499 Letter to Thomas Cooper, Nov. 29, 1802

500 Letter to Edward Carrington, January 16th, 1787

defense of “freedom of the press” did not stem from his views of “the press” at all, save his assumption that, if left free, it would produce a *diversity* of points of view. Rather than revolving centrally around his views of “the press,” Jefferson's defense of “freedom of the press” revolved primarily around his view of *the public*, particularly his modest optimism about the lay public's *judgment*. Because the lay public's judgment was sufficiently reliable that it could be counted on to recognize and remedy its errors over time, the most important feature of the informational environment was to ensure that it regularly *provided the public with the opportunity* to assess its beliefs by creating conditions likely to allow *a broad diversity of perspectives* the opportunity to reach its attention. Though Jefferson did not deny that a “press” left free to print claims at its own discretion would pepper the public with plenty of misinformation, he viewed that very same freedom as naturally providing the antidote in the form of true information – an antidote that would be recognized and taken up in the due course of time by the public, whose judgment was up to the task. This is, in an important respect, diametrically opposed to the Progressive approach, which views wide-ranging acceptance of unreliable beliefs as the natural result of an unfiltered informational environment and accordingly seeks to moderate the informational environment through a given set of authorities.

For all these reasons, I believe Madison's and Jefferson's writings reveal them to have both subscribed to something like the sort of epistemic pluralism I have described. That epistemic pluralism, I argue, lay behind the three key features of their approach I have emphasized in this chapter, which are (a) the centrality of *chronic fallibility* in their political writings; (b) their *modestly optimistic* view of lay citizens' judgment; and (c) their consequent stress on the importance of *fostering perspectival diversity* both in the informational environment and in political decision-making processes. These two thinkers' epistemic pluralism-informed approach is, I argue, decidedly superior to Progressives' institutionalism, which easily loses sight of the implications and importance of chronic fallibility, carries an excessively pessimistic view of the judgmental capacities of lay citizens, underrates the value of perspectival diversity, and is consequently incompatible with any compelling defense of liberal democracy. I believe our response to the changes to our informational environment brought on by 21st century communications technologies should be based on epistemic pluralism, should reflect trends (a), (b) and (c) in Madison and Jefferson, and should inspire us to embrace a response to those

changes by doubling down on our commitment to liberal democracy. This means rejecting the institutionalism embraced by Progressives, its pessimistic view of lay citizens, and the natural opposition to liberal democracy that come with them. I will conclude with some thoughts about how we might do this, both in terms of the public policies we embrace and in terms of the type of civic epistemology we ought to encourage in citizens and how it might be practiced in our age.

Section Two: Epistemic Pluralism and its Implications for Liberal Democracy

The sort of epistemic pluralism I have attributed to Madison and Jefferson is the best framework I have encountered for thinking about social epistemology and the growth of human knowledge in general. It is an especially good match for the sort of liberal democracy I defended in Chapter One, whose basic features of freedoms of communication and universal, equal opportunity of suffrage I believe Progressives also view as desirable. This is because the sort of epistemic pluralism I have in mind comes with the notion of *chronic* fallibility built into its very core. Yet, even as it recognizes that our fallibility is *chronic*, it sees the *active exercise of judgment* – specifically, the bringing to bear of differently-placed agents with multiple, divergent *perspectives* – as our best means of addressing this problem. And, contrary to the fears of Progressives, it does so without collapsing into the view that all perspectives are equally correct (epistemic relativism) or the view that there is no knowable truth whatsoever (epistemic nihilism). This makes it a highly useful framework for thinking about liberal democracy in the 21st century, which I will now attempt to illustrate through a series of examples.

Epistemic Pluralism's Cautious Respect for Institutions

We live in a famously large and complex political world. Dealing with a world of such size and complexity cannot be done by individuals building all their views about the world from scratch on the basis of personal experience alone. Recognition of this fact is a large part of why modern societies have formed a wide array of institutions, like those of the legacy news media and the “scientific community,” whose purpose is to preserve and advance knowledge in a given epistemic domain in ways individuals simply cannot. While these institutions are indispensable,

like all institutions they are both epistemically fallible and corruptible. Citizens of liberal democracy lose sight of these dangers at their own peril. I have argued that the institutionalism embraced by Progressives not only does not prevent this peril but actively invites it.

Epistemic pluralism, on the other hand, shores up the deficiencies in Progressives' institutionalism without requiring us to abandon the idea that institutions play a crucial role in the preservation and advancement of collective knowledge. It does so by encouraging us to see the *value* of each perspective while leaving open the possibility of giving *different weight* to different perspectives. This allows us to give more credence to certain testifiers or place more confidence in knowledge-claims based in a certain methodology than others, but it also encourages a certain degree of restraint toward even our favored sources of information and methods of inquiry by reminding us that no perspective *by itself* can lead us to the truth. This outlook allows us to resist the idea that all perspectives are equally correct while simultaneously guarding against the all-too-easy temptation to peremptorily dismiss objections to those claims we find highly persuasive. This latter tendency I have argued to be rampant in the Progressive literature objecting to the post-truth “malaise,” where the modal response to climate deniers, anti-vaccers, Brexit supporters, and the many other groups they see as the standard-bearers of the post-truth movement is to dismiss the very *possibility* that lay citizens might have legitimate objections to the claims of the experts within those expert domains. Such a peremptorily dismissive response is quite consistent with Progressives' institutionalism, which views the experts in these respective fields as possessing highly reliable methods of inquiry and the judgment of lay citizens as inadequate to the task of meaningfully critiquing the claims of these experts. However, it is *not* in line with epistemic pluralism. Instead of treating the opinions of lay citizens who resist the expert consensus on these matters as irrelevant because they lack the training and institutional credentials of the experts they oppose, epistemic pluralism would view their entry into the conversation as potentially meaningful *precisely because* they hold a different station in society, and therefore are likely to bring to bear a different, potentially insightful, point of view. Retaining our capacity to see the potential value in such critiques does not require us to abandon our beliefs about the cause and likely effects anthropogenic climate change, immunization policy, or Brexit, but it *does* encourage us to be more prepared to listen to and take lessons from the legitimate considerations being forwarded by each of these groups,

considerations which, as I argued in Chapters 2 and 3, are often kept alive by these groups and no one else. More generally, epistemic pluralism retains the ability to appreciate the pivotal role played by institutions in preserving and advancing our collective knowledge without encouraging complacency toward those institutions.

Epistemic Pluralism and Privilege

Epistemic pluralism is not just a good way of shoring up the deficiencies in the Progressive approach I have dwelt on in this dissertation. It is also better in line with certain other movements in 21st century politics viewed with favor by many Progressives. Take, for example, the ascendancy in the last several decades of conversations about *privilege* and its impact on politics and society. The central thrust of this line of argument has been to point out the ways certain citizens' upbringing and experiences in life invisibly and in many ways unavoidably impart upon them a *limited perspective* that causes them to misunderstand and/or underestimate the issues dealt with by those who do not enjoy a similar level of privilege. Many of those who forward this line of argument view increased *political representation* by people who belong to underprivileged groups as the appropriate way to address this problem. It is not clear that this view is highly compatible with the Progressive standpoint that tends to treat experts and institutionalized epistemic authorities as the primary bearers of reliable knowledge and treats lay citizens as lacking sufficient powers of judgment to meaningfully critique their claims. If anything, it seems likely to incline Progressives to dismiss the complaints of those who warn of the detrimental effects of privilege as the misguided and unwelcome rantings of a judgmentally-suspect public⁵⁰¹ and to view certain groups' preference for being represented by someone they view as *belonging to* that group, rather than someone who possesses a certain kind of expertise, as signs of anti-intellectualism.

Vaccine opposition, and mistrust of the medical community in general, is a potentially revealing example. As I showed in Chapter Three, Progressives' tendency to view the medical science community as the provisioners of truth in the domain of medicine and to view lay citizens as unfit to make legitimate judgments on matters in that domain leads them to write off

⁵⁰¹It is not uncommon to see “social justice warriors” who invoke the term “privilege” to undermine the credibility of a speaker criticized in just this way (Even 2018).

the claims of anti-vaccers, viewing them as born of ignorance, dogmatism, and other epistemic vices fostered by 21st century communications technologies. As Jessica Jaiswal and Perry N. Halkitis note, much of the public health research done on medical mistrust reflects an approach to relations between expert communities and lay citizens very similar to that of Progressives. In the words of Jaiswal and Halkitis, a “particularly substantial subset of the medical mistrust research revolves around ‘conspiracy beliefs,’” including those spread online, and views resistance to the medical community as a pathological “cultural barrier” that inhibits certain insular populations from accepting the medical consensus on issues like vaccination.⁵⁰² In the United States, such mistrust of the medical science community has long been especially prominent among Blacks. Many studies in public health attribute this mistrust primarily to the legacy of widely known historical cases of abuse of Blacks by the medical community, such as the Tuskegee experiments, which make the sorts of “conspiracy beliefs” focused on by public health researchers seem more plausible to those communities, and downplays the possibility of discrimination in the medical community itself.⁵⁰³ As Jaiswal and Halkitis point out, this Progressive-esque framing *presumes* the innocence and correctness of the contemporary medical community, as a result placing “the onus to overcome medical mistrust on [Blacks]” without apparently even considering that an equally plausible explanation for Blacks' ongoing mistrust may be that the medical science community is largely composed of “environments that engender mistrust and sustain institutionalized inequalities.”⁵⁰⁴ Given the history of race relations in the United States, surely it is not implausible that at least a significant *part* of the well-documented long-term disparity in health-related outcomes among Blacks compared to Whites is due to the performance of medical institutions themselves, and not just due to the recalcitrant refusal of Blacks to buy in to the system due to their beliefs in conspiracy theories, as seems often to be suggested in the public health literature.⁵⁰⁵ This explanation fits better with the empirical research of at least some public health scholars, such as Dwayne T. Brandon and colleagues, who found the attitudes of Blacks who knew of the Tuskegee experiments to be largely similar to

502Jaiswal and Halkitis (2019), 80-1

503E.g. Gamble (1993), Boulware et al (2003)

504Jaiswal and Halkitis (2019), 81

505Blacks' mistrust of, and consequent refusal to use, conventional medical institutions is cited among the first explanations for such disparities in numerous public health articles. See, e.g., Gamble (1993), Boulware et al (2003), Brandon et al (2005), LaVeist et al (2009), Cahill et al (2017).

those who did not, suggesting that, in the words of the authors, Blacks' mistrust of the medical science community may have less to do with their knowledge of history and the tendency of that knowledge to make them amenable to “conspiracy beliefs” and more to do with “continuing patterns of negative interactions with the healthcare system.”⁵⁰⁶

I believe this tendency among public health researchers to downplay Blacks' criticisms of medical institutions and write off their mistrust of those institutions as largely the result of “conspiracy beliefs” is an example of precisely the sort of problem warned of by those most concerned about the role of privilege in our society. It is highly likely that researchers publishing in public health journals come from comparatively privileged backgrounds that contribute to their having comparatively positive experiences with the medical science community. On top of that, public health researchers belong *to* that community themselves. Both factors are likely to skew the perspective of those researchers, likely in a way that causes them to overestimate the average performance of the medical community. This is a problem of *perspective* that cannot easily be solved by training, education, or method. Indeed, training, education and method in medical science traditions may actively *contribute* to the problem by making the researchers more understanding of and sympathetic to the medical science point-of-view. Progressives' selectively fallible version of institutionalism, which views this sort of training and adoption of established methods as the *very reason that* institutions like the medical science community produce more reliable knowledge than lay citizens, as well as why lay citizens are largely unfit to critique those institutions, offers no solution to these problems. On the contrary, it exacerbates them. Public health researchers' tendency not to seriously consider the complicity of the medical science community in the health outcomes of Blacks and to write off Blacks' mistrust of the medical community as stemming mainly from “conspiracy beliefs” encouraged by their over-attention to sensational incidents like the Tuskegee Study reflects a view of medical scientists as either infallible in the domain of medical science or at least so reliable that we need not even seriously consider that their claims and practices may be either corrupt or seriously in error. At the same time, it reflects a pessimistic view of lay citizens resistant to that community by heavily emphasizing their deficiencies in judgment (e.g. “conspiracy thinking”) as explanations for their mistrust. In other words, it reflects precisely the same combination of veneration toward favored

⁵⁰⁶Brandon et al (2005)

epistemic authorities and disparagement of the judgment of lay citizens as that characteristic of Progressives. Far from helping us solve the problems of privilege, Progressivism creates an enabling environment for those problems to continue to plague medical science research and all the other avenues of life in which power-holders are likely to hold a different perspective from those who do not hold it.

Where Progressivism creates an enabling environment for these problems, epistemic pluralism actively combats them. It does so by encouraging us to view *both* Blacks outside the medical science community *and* public health researchers as evaluating the *same world* from a *different perspective*. Because they world they are seeing is the same but each perspective only sees it from a certain angle, epistemic pluralism expects each perspective *both* to hold potentially valuable insights *and* to contain some blind spots. Public health researchers are likely to have a better macro-level view of the correlations between certain health programs, such as vaccines, and outcomes, such as the instance of vaccine-related illness. They are *not* likely, however, to have an accurate view of the on-the-ground conditions that face individual patients at local medical facilities. They are not likely to have a good sense – or indeed any sense at all – of the level of attentiveness, prejudice, integrity, sincerity, propensity to care, or any other of the numerous qualitative features of health care provision that affect the quality of the health care given to patients on the ground by doctors, nurses, secretaries and administrators. These are not irrelevant variables but are instead data of fundamental importance in determining the extent to which the pervasively different health care outcomes experienced by Blacks in the contemporary United States are due to widespread prejudice, as opposed to being attributable to the behaviors and attitudes of Blacks themselves or other factors. Among the people best placed to make these assessments are Black lay citizens themselves. Their *perspective* is valuable, indeed *indispensable*, in determining the full truth about the American health care system. Progressivism's selectively skeptical version of institutionalism seems to encourage us to presume the innocence and correctness of the medical science community, or at least not to take the danger of their making mistakes and/or being corrupt very seriously, and to view the mistrust of Blacks as likely caused by their deficiencies in judgment. Meanwhile, epistemic pluralism encourages us to seek out the *value* in *each* perspective, determine their appropriate weight, and triangulate the truth it expects to lie somewhere in-between.

Epistemic Pluralism's Affinity with Liberal Democracy

These same considerations can be generalized to demonstrate why this sort of epistemic pluralism is an especially attractive match for both liberalism and democracy and, *a fortiori*, liberal democracy, as well as why Progressives' selectively fallible institutionalism is not. Despite Progressives' apparent self-identification as liberal democrats, the Progressive standpoint's heavy emphasis on the benefits of *expertise* and routine deference to authority and its concurrent emphasis on the judgmental deficiencies of *lay citizens* make it a very odd match for liberal democracy. As I argued in Chapter One, a coherent (but not necessarily compelling) case could be made by Progressives for liberalism *without* democracy, viewing the distorted beliefs likely to result from the interaction of liberal freedoms with citizens' deeply flawed judgment as largely harmless since important political decisions are not left up to lay citizens.⁵⁰⁷ Likewise, the Progressive view is compatible with a defense of democracy *without* liberalism, since authorities would be able to control the information available to lay citizens, thus reducing the likelihood they will be deceived by misinformation and demagoguery.⁵⁰⁸ The combination of both liberalism *and* democracy, however, does not make sense. Indeed, it seems fundamentally objectionable, since Progressives clearly expect the result of liberal freedoms of communication let loose upon a public possessed of such impoverished powers of judgment to be a rise in the influence of demagogues and other opinion manipulators and the general reign of misinformation on the informational environment and in the beliefs of the public. Why would anyone who held such a view of the public while simultaneously holding a view of certain institutions as highly reliable and trustworthy advocate both leaving the informational environment unmanaged by authorities in the way championed by liberals *and* leaving political power in the hands of such an easily misled public in the way championed by democrats? It seems far more reasonable for someone subscribing to an institutionalism like that favored by Progressives' to prefer to place control in the hands of the institutions they hold in such high regard, either by allowing them to control the informational environment or by taking decision-making power out of the hands of the public and putting it into the surer hands of these

⁵⁰⁷This is the essence of Brennan's "epistocracy" (2016).

⁵⁰⁸I take this to be the modal Progressive position, though most Progressives do not acknowledge it as distinctly illiberal.

institutions. Liberal democracy is, in short, a poor match for Progressives' version of institutionalism.

Epistemic pluralism, in contrast, is a natural fit for both liberalism *and* democracy and, *a fortiori*, liberal democracy. Epistemic pluralism is a match for liberalism because it offers an intuitive defense of liberalism's freedoms of communication. Since all perspectives are simultaneously *fallible* yet likely to hold at least *some* potential insight into the world we share in common, it makes sense to endorse a policy emphasizing the flourishing of a *wide variety of perspectives*, rather than trying to identify the “best” perspectives and arrange things so they predominate. I take this to be a large part of the canonical defense of liberal freedoms of communication found in Chapter Two of John Stuart Mill's *On Liberty*. Epistemic pluralism is also a match for democracy because it makes sense of our defense of universal, equal opportunity of suffrage among able-minded adults, rather than opting for technocracy or some other form of “epistocracy.” Because each of the adults casting a vote in a democracy is differently situated, each is likely to have a unique, and likely valuable, view of the world we share in common, the problems that beset us as a community, and the solutions that make sense to address those problems.⁵⁰⁹ While the perspectives of experts may have special value in identifying the nature and best solutions to some of those problems, relying solely on that perspective cannot yield a complete a view of the whole. Instead, it must be supplemented by the views of lay persons, whose value comes precisely from the fact that they are *not* experts, instead occupying a standpoint that allows them to see the world they share in common from a different point of view, offering potentially valuable insights to which experts, by virtue of their position as experts, *cannot* be privy. When we view things this way, universal, equal opportunity of suffrage makes sense because it seeks to recruit *all* perspectives into the public decision-making process, thus giving society as a whole the maximum array of points-of-view through which to try to triangulate the truth.

Epistemic pluralism is as compatible with liberalism and democracy in combination as it each with each individually. Whereas the combination of liberalism and democracy seems to be the worst of all worlds under Progressives' institutionalism, epistemic pluralism views them as possessing a synergistic relationship. On this view, liberal freedoms of communication provide

⁵⁰⁹This line of argument has been well prosecuted by Landemore (2012) and Page (2007).

an *epistemic* mechanism by which citizens may come to steadily improve their understanding of the world they share in common – which includes improving their beliefs generally as well as becoming apprised of possible instances of corruption – by hearing what the world looks like from the perspective of a wide variety of other, differently-situated citizens and groups. Democracy's universal, equal opportunity of suffrage, meanwhile, provides a *practical* mechanism for infusing perspectival diversity into political decision-making processes. Since citizens' beliefs are improving and their ability to identify illegitimate authority is maximized, citizens should converge – in the long run and over time – upon similar, or at least overlapping, views of the world, allowing them to identify common problems and eventually agree upon solutions to them. Were liberalism *not* to be combined with democracy, however, these expected benefits would be greatly truncated. Liberalism in the absence of democracy would result in public policies being informed by too narrow a range of perspectives, making it more likely they would be short-sighted and/or characterized by corruption. Democracy in the absence of liberalism would deprive citizens of the best-possible conditions for improving their beliefs and discovering illegitimate authority, thus hampering them from converging on common solutions to common problems and encouraging the development of corruption. In these ways, epistemic pluralism reveals an affinity to liberal democracy that sharply contrasts with the implications of Progressives' institutionalism.

Epistemic Pluralism's View of the Value of Scientific and Journalistic Best Practices

My emphasis so far might give the impression that epistemic pluralism is sympathetic with lay citizens and “outsiders” and antithetical to established institutions like the “scientific community” and the legacy news media. However, far from being *antithetical* to such institutions, I believe epistemic pluralism offers the *best explanation for* why, when things are working well, institutions like these have special epistemic value to society. Recall the arguments of Goldman in defense of these two institutions, related in Chapter Five. In that argument, Goldman mentioned peer-review as a practice employed by scientists that made the claims of the “scientific community” more reliable, on average, than the claims found on the blogosphere. While I agree with Goldman that there is great epistemic value in peer-review, I believe that he misidentified where that value comes from. Goldman's argument was that this demonstrated the

epistemic benefits of information *filtering by the relevant authorities*.⁵¹⁰ However, I believe the primary value of peer-review does come not from the fact that it is an instance of filtering by the relevant authority, but from the fact that it is an institutionalized mechanism for *bringing perspectival diversity* to bear on a subject. The same is true of the other feature of science that has most contributed to the high degree of reliability of scientific claims over time, which is the practice of *replicability*. The ideal standard is for peer-review be both “blind” and performed by multiple, unrelated reviewers. The latter requirement is a tacit admission that no single authority possesses sufficiently reliable judgment and character to deserve to be treated as the ultimate arbiter of quality in any scientific domain. As for replicability, it obviously does not suffice for the replication of a given study to be performed and reported by the *same* researcher or research team as was responsible for the original finding. Instead, it must come from somewhere else. As long as they carry out the replication competently, the *more different* that “someone” can be from the original researcher – where by “more different” I mean divergent from them in terms of the likely interests at stake and in terms of ideological affinity – the better. If scientists funded by Exxon-Mobile arrive at the same findings as climatologists funded by the Sierra Club, that finding seems more weighty than if the only scientists replicating the Sierra Club-funded team's results are other teams funded by organizations with a similar political agenda. The gold standard of replicability is for *everyone* who attempts the experiment to get the *same* results no matter what their antecedent point-of-view.

The two main things peer-review and replicability are meant to combat are *corruption* and *good-faith errors*. Peer-review is supposed to be anonymous to try and ensure that reviewers do not have a conflict of interest that might either distort their judgment or inspire them to commend or reject articles on an ideological or self-interested basis. Replicability is ideally performed by practitioners unrelated to the original researchers for the same reasons. Meanwhile, both peer-review and replicability are meant to catch good-faith errors, such as flaws in the original researchers' reasoning, methodological design, or execution in performing the study. In other words, *the best science is science that has been vetted by a diversity of perspectives*. It is *this* perspectival diversity, and not the mere imposition of filtering by authority emphasized by Goldman, that best contributes to the reliability of scientific research.

⁵¹⁰Goldman (2009), 6

A very similar argument applies to Goldman's defense of another favored Progressive institution, that of the legacy news media. As shown in Chapter Five, Goldman defends the reliability of the legacy news media on the same grounds as form the basis of his defense of the “scientific community.” However, I believe each of the specific practices he cites as contributing to the epistemic reliability of the legacy news media draws its epistemic value not chiefly from the mere fact that some sort of authority is imposed between the process of information-gathering and publication but instead from its capacity to facilitate the sort of perspectival triangulation described by Arendt and Conant. Recall that the practices Goldman cites are the employment of “fact-checkers,” the requirement for journalists to consult more than one source before publishing an article, and industry-standard limitations on the use of anonymous sources.⁵¹¹ The first two of these are clearly instances of injecting perspectival diversity into the pre-publication vetting process. Though less obvious, perspectival considerations are also behind the best defenses of the need to limit use of anonymous sources. The demand that journalists limit their reliance on unnamed sources is justified primarily due to the way the use of such sources prevents readers from being able to *evaluate the perspective* of the informant(s) being relied upon for a journalist's stories and condition their belief on that evaluation's basis, and so inhibits them from being able to triangulate to the truth by anticipating the likely bias of that source. The reasons these perspectivism-driven standards are endorsed are the same as the reasons the standards of peer-review and replicability are endorsed by the scientific community. They are meant to combat the dangers of corruption and good-faith errors. Again, the primary value of these practices is not that they are mechanisms of filtering by authority, but instead their role in attempting to ensure that perspectival diversity is brought to bear on the articles that go through the publications process.

Epistemic pluralism is not a sweeping vindication of the views of lay citizens and “outsiders” because it recognizes that they, like everyone else, are fallible and limited. Nor is it a sweeping condemnation of knowledge-producing institutions like the “scientific community” and the legacy news media. On the contrary, it offers a compelling framework for explaining why we might expect institutions like these to produce especially reliable knowledge (so long as they are working as intended), a framework that, unlike the selectively fallible institutionalism of

⁵¹¹Ibid., 7

Progressives, keeps clearly in view the need to combat those institutions' members' moral and epistemic fallibility.

The Role of Judgment in Epistemic Pluralism

Though I put it in parentheses above, the “if they are working as intended” part is of crucial importance. It is not something that can simply be assumed by citizens. Instead, it is something they themselves must assess. This requirement brings up a faculty I have not yet highlighted in this chapter, yet which is integral both to my argument in this dissertation as a whole and to the account of epistemic pluralism I am defending in this chapter. This is the faculty of *judgment*. The sort of epistemic pluralism I am defending here sees the active and assertive exercise of judgment – what I have called doxastic self-rule – as a *crucial aspect of and contributor to* our collective well-being. Accordingly, any model of civic epistemology based on it will see the regular practice of doxastic self-rule by citizens as an asset to democracy, rather than a threat to its well-being.

As we have seen, Progressives often view the judgment of lay citizens not as an asset but as a liability. As I argued in Chapters Two and Four, this view is behind the pervasive Progressive view that by decreasing the barriers to publication, 21st century communications technologies as problematic for democracy, a view that only makes sense if we assume lay citizens' judgment is so poor that they will be largely unable to differentiate truth from falsehood encountered online. I have already disputed this argument as incompatible with any convincing defense of liberal democracy (Chapter One) as well as on empirical grounds, arguing that the idea that decreased barriers to publication is likely to result in decreased quality in lay citizens' beliefs is strongly contradicted by historical trends (Chapter Two) and that the supposition that citizens' judgment cannot be relied upon to differentiate between truth and falsehood online is strongly contradicted by empirical research on fake news (Chapter Four). Against Progressives' deep pessimism about the judgment of lay citizens and their consequent embrace of a deferential model of civic epistemology, I have defended a civic epistemology based a certain *modest optimism* about lay citizens' capacities of judgment that encourages them to assertively exercise *doxastic self-rule*, insisting that their beliefs and decisions about whom to trust be based on their *own judgment*, and not the result of routine deference to pre-defined authorities.

Epistemic pluralism embodies just this sort of modest optimism and justifies just this sort of doxastic self-rule. Its *optimism* about lay citizens' judgment is implied by its assumption that each perspective is a view of the "world we share in common" and is likely to bear some meaningful, unique insight into that world. Yet, that optimism is *modest* because it also recognizes that all perspectives see that world from only one point-of-view, and therefore each one is *limited* and *fallible*. While recognizing that every individual citizen's perspective is limited and fallible, epistemic pluralism nevertheless implies support for the *assertive* exercise of doxastic self-rule by those citizens. This is because the value those individual citizens bring to the community *is* their perspective. It is precisely the uniqueness and idiosyncrasy *of* their perspective – the same feature that means it is inevitably bound to be flawed – that lends it the potential to offer insight about the world we share. epistemic pluralism encourages individuals to exercise doxastic self-rule because it recognizes that, flawed though they may be, individuals' judgments about the real world are just that – judgments *about the real world*, the world we share in common. On this view, citizens' first task is to *see* the world the way it *appears to them*. They then contribute to the community as a whole by offering *their* views up for everyone else to consider. When each citizen does this, the collective is offered access to a wealth of different views – a panoply of perspectives to weigh and consider and, by following the trajectory of those they find most persuasive, attempt to triangulate a better sense of the shape of the world we share in common than any individual perspective might form on its own. epistemic pluralism sees the view of the world likely to emerge from this process of triangulation as *more accurate* than the view of the world any of us could reach by consulting a single perspective, whether that be the perspective of an individual left to their own epistemic resources or the perspective of a single research tradition, methodology, or knowledge-producing institution.

This language may seem a bit starry-eyed and abstract, but I believe the examples given in the previous sub-sections show epistemic pluralism to have great practical value and carry important implications for the real world. Epistemic pluralism offers an elegant framework for making sense of contemporary groups' arguments about the pervasive and insidious dangers posed by privilege. It also explains why some well-designed institutions may reasonably be expected to produce and disseminate more consistently reliable knowledge than individuals by themselves or institutions that are less well-designed, while simultaneously encouraging citizens

to engage in the oversight indispensable to making sure those institutions continue to perform well. In both these cases, epistemic pluralism sees the *active exercise of judgment* by multiple individuals bearing multiple different perspectives as the very *means by which* corruption is most effectively guarded against and collective knowledge is most reliably improved. Its approach to fighting corruption and improving our collective epistemic well-being is, therefore, not to reserve the task of knowledge-seeking solely or primarily to some specialized group, demanding that other groups routinely defer to that group's claims and condemning those who refuse as threats to the public good. In many important respects it is the opposite: epistemic pluralism seeks to foster the growth of collective knowledge and guard against the possibility of corruption by encouraging differently-placed individuals to come to their *own* conclusions about what to believe and whom to trust and to *air those views publicly*, thus granting the public as a whole a broad spectrum of points-of-view through which to view the world they share in common – more angles from which to look at that world in an attempt to triangulate the truth.

Epistemic Pluralism's Opposition to Censorship on the Basis of Fact

I believe an embrace of epistemic pluralism has a few clear implications for public policy and civic epistemology. I will begin with the former. The most obvious public policy implication I want to emphasize here has to do with how we manage – or, preferably, decline to manage – information online. As I showed in Chapter Two, scholars in a wide array of disciplines have for the last decade been preoccupied with the topic of misinformation online, which many view to be so damaging as to justify its being categorized as a public health issue.⁵¹² Yet, as I argued in Chapter Two, scholarly treatment of this issue suffers from a number of methodological and conceptual shortcomings. These shortcomings all share in common the *presupposition* of a deeply pessimistic view of lay citizens' judgment by the commentators, which biases their analyses in ways that lead them to adopt premature, and perhaps fundamentally misguided, conclusions about the plight of our collective epistemic well-being. In Chapter Two I argued that

⁵¹² This is the implication of Hossein Derakhshan and Christine Wardle's popular framing of contemporary citizens as widely afflicted by a “misinformation disorder” (Wardle and Derakhshan 2017). Indeed, so many commentators use medical terms to describe the epistemic state of 21st century democracy due to misinformation online that Johan Farkas and Jannick Schou call medical analogies “the discursive backbone of the democratic imaginary currently conjured up in scholarly, journalistic, and political discourses,” citing dozens of examples (Jannick and Schou 2020, 47-49).

it is not at all clear that either historical trends or the best available empirical evidence support such pessimistic presuppositions or conclusions, an argument further supported by the empirical research on “fake news” cited at the end of Chapter Four. Similarly dubious is the worry fairly commonly voiced by Progressives that citizens' growing tendency to resist the claims of Progressives' favored epistemic authorities is evidence of those citizens descending into epistemic relativism, mentioned above in Section One. Very few people outside intellectual or mystical circles have ever genuinely subscribed to such a view, which is so extremely impractical and so extraordinarily far removed from common sense thinking that I doubt it will ever be found attractive outside those confines.

It is in response to these thinly supported, and perhaps fundamentally misguided, worries that Progressives argue in favor of imposing some sort of authority over the informational environment. Often enough, these recommendations come with a certain amount of hand-wringing as Progressives recognize the dissonance between their avowed support for liberal democracy and their preference for the imposition of some form of authority on the informational environment. Anya Schiffrin, for instance, expresses the desire to avoid “corporate and government censorship,” but follows this up by implying that the former is precisely the policy she supports, since “it is important not to let technology companies use free speech as an excuse not to take action.”⁵¹³ Like many Progressives, Schiffrin seems driven to sanction such anti-liberal policies by her perception that 21st century communications technologies have thrown us into a state of collective epistemic emergency. Yet, Schiffrin's justification of this sense of emergency is strikingly thin. Speaking of her worries about the much-touted Russian interference in the 2016 American Presidential election race, Schiffrin asserts that we “simply do not know what it means for the electorate when millions of Russian propaganda messages are targeted at swing states.” Yet, she follows it up by immediately saying that “Even so, it is not too early to take action. When there is a strong possibility of danger, society must act.”⁵¹⁴ Similarly, Nathaniel Persily in one paragraph admits that we “do not yet know how big an effect fake news had on the 2016 campaign,” yet proceeds several paragraphs later to write of the profound “power of fake news,” which “create[s] a blanket of fog that obscures the real news” and whose

⁵¹³ Schiffrin (2017), 122

⁵¹⁴ Ibid., 121

“prevalence...online erects barriers to educated political decision making and renders it less likely that voters will choose on the basis of genuine information rather than lies or misleading 'spin.’”⁵¹⁵ These inspire Persily to emphasize the importance to democracy of “strong intermediary institutions that can act as guardrails to channel decision-making.”⁵¹⁶

What stands out most glaringly in these two accounts, and in Progressive accounts more generally, is their rush to embrace institutional authorities even after admitting they lack anything resembling strong empirical grounds for coming to deeply pessimistic conclusions about the actual effects of 21st century communications technologies on the beliefs and behaviors of lay citizens. Since that impulse cannot be explained by the vividness of the empirical evidence, it must stem from something else. I believe that “something” is their embrace of Progressives' institutionalism and its combination of pessimism about the judgmental capacities of lay citizens and uncritical veneration of certain knowledge-producing institutions. The depth of Schiffrin's and Persily's pessimism about the judgment of lay citizens is evident in their mutual assumption – an assumption universal among Progressives – that citizens are largely helpless to sort out fact from fiction in the face of the volume of misinformation they encounter online, an assumption so strong that it impels them to make confident assertions about the detrimental epistemic impact of 21st century communications technologies even *after* explicitly admitting their own lack of evidence for those conclusions. It is the combination of this deep pessimism about lay citizens with an insufficiently critical view of the favored knowledge-producing institutions that impels them to so strongly emphasize the democratic importance of epistemic authorities and to endorse the need for them to serve as mediators in the informational environment with only a token gesture toward the possible dangers of doing so.

I argue against both this way of thinking about the judgment of lay citizens and this way of conceptualizing the role of institutional mediators in the informational environment. There is no good reason to suppose that 21st century communications technologies will overwhelm citizens' general ability to tell truth from fiction, any more than the countless other technological developments that have continually expanded the range of information available to citizens since the invention of the printing press. To the contrary, the historical record and current empirical

515 Persily (2017), 69-70

516 Ibid., 72

evidence cited throughout this dissertation would seem to suggest a rather resounding vindication of the optimism of Jefferson, who trusted that the “light which has been shed on mankind by the art of printing has eminently changed the condition of the world” for the better, to the degree that “while printing is preserved, it can no more recede than the sun return on his course.”⁵¹⁷ While there is no good reason to suppose the internet will overwhelm citizens' cognitive capacities any more than the numerous advances in communications technologies that similarly worried intellectuals in the past, there *is* good reason to suppose that any institution granted the power to censor information online will not only make mistakes but will likely use that power to further its own interests at the expense of the public. That is the characteristic pattern of political regimes in all contemporary countries that do not protect freedom of speech and is the consistent historical pattern of regimes in prior eras. I see no reason why we should expect any different behavior from any non-governmental authority given the same power and many reasons why we should expect the same. Since there would seem to be good reason to hold a *modestly optimistic* view of the judgment of the lay public and abundant reason to hold a *deeply pessimistic* view of the likely use to which any institution given the power would censor people's information, I argue there is no good reason to follow Progressives in advocating censorship either by “fact-checkers” at Facebook, Google, or Twitter or by government authorities. There is no good reason to suppose *either* that citizens are so unequipped to deal with misinformation that it poses a threat to democracy *or* that the imposition of any authority in a position of censorship is likely to *improve* the epistemic condition of the informational environment. If anything, the best-available evidence would suggest just the opposite on both counts.

Epistemic Pluralism's Preference for Idiosyncratic, not Uniform, Information Filtering

This does not, however, suggest that the information made available to citizens' online need not be filtered *at all*. On the contrary, as Cass Sunstein rightly points out, information filtering is “inevitable in order to avoid [information] overload and impose some order on an overwhelming number of sources of information.”⁵¹⁸ As Sunstein explains, the sorts of

⁵¹⁷ Letter to John Adams, September 4th, 1823

⁵¹⁸ Sunstein (2016), 64

algorithms that have become the subject of heated debate have become part of our daily lives not *only* because they help tech companies hold your attention and market your information to advertisers (although that is a big part of *how* they are designed) but also out of the *sheer necessity* of sorting through the unimaginable ocean of information that exists online. Filtering is an essential, and often extremely beneficial, aspect of what companies like Google, Facebook and Twitter do. It is not filtering *per se* we should oppose but instead the *uniform* filtering of *most or all citizens'* information by the *same* authorities or according to the *same* standards. Consider the case of political shock-jock Alex Jones and his organization *Infowars*. Epistemic pluralism need not object to Jones being removed from one or even several platforms, such as Facebook and Youtube, *so long as* that removal was the result of those platforms having *independently* reached the decision to ban that content on the grounds of its having violated that platform's *own unique* set of standards. So long as each individual platform is independently forming its *own* set of standards and employing them according to its *own* authorities' independently exercised judgment, the overall informational environment should be populated by a wide diversity of views for citizens to assess in the light of their own judgment. I call this sort of filtering by individual institutions, each with its own set of standards based on its own priorities, perspective and vision, *idiosyncratic* information filtering. It is *idiosyncratic* because it should not be expected to put a *uniform* filter on *all*, or practically all, citizens' information and should not result in a *homogenous* epistemic environment.

Another example of idiosyncratic filtering is the sort of filtering that results from the development of unique algorithms based on each individual internet user's web behavior. Since each individual's decisions are unique, each individual user's algorithm is different, meaning that the content to which they are exposed on their Facebook feed, Google results, etc., is unique. Likely, this would encourage citizens to form different views of the world – different perspectives. The epistemic pluralist would see no problem with this. In fact, she would likely view that development as desirable, as it would mean that many different perspectives were being kept alive within society as a whole. The epistemic pluralist would simply want to make sure that such perspectives were allowed, or even encouraged, to interact with one another frequently enough to make sure that significant segments of the public are able to be exposed the perspectives that develop partly through this sort of algorithmic sorting. Empirical research

suggests there is every reason for optimism on both fronts. People are almost certainly exposed to a far wider array of views online than in their personal lives⁵¹⁹ and most internet users' online environments result in their being exposed to a significant degree of ideological diversity.⁵²⁰ The epistemic pluralist would view this as a highly promising state of affairs.

This is in many ways the *opposite* of the approach taken by Progressives. Progressives tend to view *idiosyncratic* filtering as problematic for democracy and to advocate *uniform* filtering in its place. Rather than encouraging different online platforms to develop and enforce their own policies and standards, Progressives encourage *all* platforms to *uniformly* filter out views they find “dangerous” to democracy. Inevitably, these views are characterized as “conspiracy beliefs” or summarized in other terms whose real meaning, on close examination, simply communicates their estimated implausibility or offensiveness in the eyes of the commentator and their anticipated readership. Disturbingly, this type of censorship appears to be on the rise. Jones and *Infowars* is a prime example. Rather than being removed by just one or a couple companies at different points in time based on the judgments of the relevant authorities, Jones' pages were suddenly removed from *all* major platforms, including Facebook, Twitter and Youtube, within the space of a few days. Shortly afterward, Jones' web domain, “Infowars,” was revoked, after which other domain providers refused to host his content. It seems obvious that Jones' removal was not the result of each of these companies exerting its own judgment and coming to the conclusion that Jones' content violated their standards, but was instead the result of all these companies either following the lead of the first mover or responding to the *same* or *similar* public pressure campaigns. Whichever the case, the end result is the same: the uniform, sweeping, indeed nearly comprehensive silencing of a voice from the informational environment in its entirety. Epistemic pluralists would view this as a troubling precedent, as they would expect this sort of silencing, if normalized, to inevitably lead to both a decreased collective ability to improve our knowledge and an increased ability for authorities – especially those imposing the censoring – to engage in corruption without fear of discovery.

Another example can be found in major platforms' response to the recent COVID-19 pandemic, during which Facebook, Google, Twitter and other platforms all made intentional

519Gentzkow and Shapiro (2011)

520Flaxman et al (2016), Garrett (2017)

changes to their algorithms so that U.S. users searching for information about COVID-19 would be directed to the *same* sources of information, all of which were U.S. government agencies. Health officials and numerous public commentators “welcomed the new effort,” presuming it to help “scrub [those] sites of coronavirus misinformation.”⁵²¹ Support for this information-filtering effort came on the back of the same combination of deep pessimism about the lay public's judgment and uncritical faith in the reliability of large institutions I have argued to be characteristic of Progressives. Commentators who embraced such efforts on the part of Google, Facebook and Twitter were preoccupied with the threat of large numbers of the public embracing misinformation about COVID-19 but apparently unworried about the possibility that major online platforms and the U.S. government might *themselves* be imperfect, or even corrupt, information providers. Some readers may be inclined to write off the possibility of the latter as “conspiracy thinking.” I believe such thinking is naive, however, for several reasons. First, the emergent nature of COVID-19 meant *no* authority could credibly claim to understand it sufficiently to determine the “truth” about the virus at the time the platforms put these filtering measures into place in mid-2020. Second, there is no good reason to suppose either large corporations or the government are likely to decide what information to release about COVID-19 or any other issue purely on the basis of the public interest. Third, in the very case at hand the chief medical advisor to the president, Dr. Anthony Fauci, publicly admitted to providing misinformation to the public by claiming in the early period of the outbreak that wearing masks was more likely to spread COVID-19 than prevent it. Why should the American public suppose this was the *only* way in which such information was skewed or distorted by the authorities, especially given that Fauci has publicly stated that he doesn't “regret anything” about the way he handled the mask issue?⁵²² Epistemic pluralism does not suppose *that* either these corporations or the U.S. government systematically lied to the public about COVID-19 or that the information filtered out of users' search results was the truth, but it *does* oppose both the extraordinary pessimism about the judgment of the public implicit in commentators' assumption that misinformation about COVID-19 was bound to lead to massive public harms as well as the easy, uncritical assumption that the filtering done by major platforms was done accurately and

521 “How Google, Facebook and Others are Using Algorithms” (2020)

522 O'Donnell (2020)

benevolently and that the same information provided by the U.S. agencies was similarly accurate and benignly motivated. Accordingly, epistemic pluralism would be far more hesitant to embrace such filtering, and far more critical of the capacities and motives of those doing the filtering, than the public health officials who readily endorsed it.

The reasoning behind these epistemic pluralist positions reflects the epistemic pluralist's appreciation of chronic fallibilism and the consequent need for diversity. Subjecting the information available to *most or all* citizens to the *same* filtering methodology *imposes a uniformity of perspective* on once-disparate contents. Whatever makes it through the filter is what a *single or narrow range of perspective(s)* – the perspective of an individual, group, methodology, or narrow range thereof – deems to be of value. The narrower this range of perspectives – the more uniform the set of authorities doing the filtering or the process they use to filter the information left to citizens – the more citizens' ability to discover new truths and uncover illegitimate authority is hampered. This is because discovery depends, as I have argued, heavily on the availability of multiple perspectives, *especially* those that challenge prevailing beliefs and established authorities. These perspectives are, of course, the very perspectives most likely to be opposed by the authorities.

While epistemic pluralism opposes the kind of *sweeping, uniform* filtering of information in the examples given in the previous paragraph, it does *not* oppose what I have called *idiosyncratic* information filtering. Idiosyncratic information filtering *does* involve limiting individual users' exposure to certain kinds of information. However, it still allows the proliferation of *multiple perspectives* to proliferate in the informational environment as a whole. Undoubtedly this allows many inaccurate views to persist and sometimes even thrive among the public. However, this does not worry the epistemic pluralist excessively because the epistemic pluralist is not as concerned about the accuracy of the contents of any particular user's “feed” or the excellence of any particular platform's terms of use as about the *diversity* of information allowed to flourish in the informational environment as a whole. It views an informational environment that facilitates the flourishing of a wide variety of content, determined by a multitude of unique algorithms and curated through different platform-determined criteria, as a sign of the *robustness* of the informational environment – an informational environment that gives citizens the tools they need to contribute to and benefit from the continual growth of

knowledge and the ongoing monitoring of power holders over time. Moreover, it views the continued existence in that informational environment of groups subscribing to views held by power holders and the mainstream as outlandish, “extreme” and even “dangerous” as a *positive* sign, for without them we cannot be sure our informational environment is facilitating the sorts of innovative and/or critical points of view necessary to the growth of knowledge and uncovering of corruption. It is able to do so with confidence because of its *modest optimism* about the public's judgment, an optimism it views as largely vindicated by the history of liberal democracy and by the best-available empirical indicators at present.

Again, this view is diametrically opposed to that of Progressives, who favor *sweeping, uniform* information filtering and view *idiosyncratic* filtering as a threat to democracy. The tailoring of algorithms to individual users is, for example, mourned by Progressives as the source of “echo chambers” and “filter bubbles,” which they in turn connect to a host of other societal ills while speaking highly of the broadcast era of media, where citizens' news of the world was mediated by a few, largely uniform legacy news institutions.⁵²³ Subsequent research strongly suggests these worries, like Progressive worries about fake news, are largely overstated.⁵²⁴ Even were this not the case, however, an implication of my arguments in Chapter Three regarding how local-level vices can actually be collective-level virtues suggests that an epistemic pluralist would not see the *mere* existence, or even proliferation, of “echo chambers” and “filter bubbles” as necessarily a bad sign for democracy. While an epistemic pluralist would acknowledge that *some* echo chambers and filter bubbles are likely to be repositories for deeply misguided and destructive ideas, they would be inclined to view *most* such communities as laboratories for the development and dissemination of *flawed but informative perspectives*. The epistemic pluralist would therefore see the sorting of citizens into these micro-communities as a potentially valuable *asset* to democracy and therefore be largely inclined to *encourage* their development. Indeed, as suggested in Chapter Three, I believe this very phenomenon to have already benefited democracy by allowing historically disparaged communities, such as those who share a “deviant” sexuality, to develop solidarity and grow into an influential political force.

⁵²³ See, e.g., Pariser (2011) and Sunstein (2016)

⁵²⁴ See Garrett (2017) for citations of the most influential social science research on this topic.

The epistemic pluralist would note that Progressives' approval for certain kinds of *uniform* information filtering always reduces to Progressives' approval of silencing ideas and perspectives *they themselves* find inaccurate, abhorrent, or otherwise “damaging” to democracy, either on a case-by-case basis or by the sanctioning of an authority they trust – often enough without much explanation – to do the job accurately and honestly. In addition to disagreeing with the notion that lay citizens' judgment is so impoverished as to *require* such uniform filtering, perspectivists in the tradition of Madison and Jefferson would be appalled at the evident lack of any notion of chronic fallibility inherent in such a position, as well as the ease with which it would invite corruption. So should 21st century democratic citizens. They should therefore oppose such uniform filtering in all its forms and be more sympathetic to the potential benefits of idiosyncratic filtering, even if – or, more accurately, *largely because* – it results in “echo chambers” and “filter bubbles.” Those “echo chambers” and “filter bubbles” are potentially invaluable repositories of innovation and critique, facilitating the proliferation of precisely the sorts of unique and idiosyncratic perspectives democracy needs in order to flourish epistemically.

Implications of Epistemic Pluralism for Civic Epistemology

Just as 21st century democratic citizens should oppose the forms of sweeping, universal information filtering supported by Progressives, they should oppose their selectively skeptical and deferential model of civic epistemology. That model is predicated on the same institutionalism that underwrites Progressives' embrace of universal filtering and, if adopted, would lead to the same problems. Specifically, the adoption by citizens of that model's insufficiently critical attitude toward established epistemic authorities along with its aversion to doxastic self-rule would invite complacency among those authorities, thus diminishing their epistemic value over time and actively inviting the development of corruption.

This can be shown by looking at Progressives' attitudes toward the two institutions that have served as running examples of the Progressive approach throughout this dissertation, the “scientific community” and the legacy news media. In Chapter Three, I noted that Progressives frequently chastise anti-vaccers and portray them as examples of the democratically threatening nature of the post-truth “malaise.” In that same chapter, I noted that these same Progressives

almost unanimously ignore the question raised by prominent anti-vaccination groups of why vaccine producers in the United States are not eligible to be held legally liable if something goes wrong with their vaccines, despite that question being provocative enough to be *interesting*, at the very least, and of obvious public importance. Later on, I related the saga of the *Lancet* article that used fabricated evidence to falsely pose the connection between autism and the MMR vaccine now frequently cited by anti-vaccers, arguing that although Progressives tend to treat the “MMR vaccine causes autism” claim as *anti-scientific* and the retraction of the original *Lancet* article as an example of “science policing science,” in reality that myth *began* in the “scientific community” and was only *dispelled* by after the determined efforts of lay citizens, and even then only after much resistance by the *Lancet*. Similarly, Progressives' characterizations of anti-vaccers as irrational loonies indicates that they do not take seriously the real, extensive and in some cases shockingly recent history of medical abuses to peoples of color around the world under the guise of vaccine administration.

I believe all these examples illustrate the *insufficiently critical* approach to established epistemic authorities invited by Progressives' institutionalism. That approach motivates Progressives to seek out and notice fallibility only as it manifests in lay citizens, such as anti-vaccers, and to overlook it when it manifests in the institutions they view with favor, such as the scientists who publish in and edit *The Lancet*. Were citizens to universally adopt such an attitude, it is not clear that the sort of criticism that finally led to the retraction of the article falsely attributed the onset of autism to the administration of the MMR vaccine would arise, or that the abuses administered by numerous regimes under the guise of public health would have been uncovered. Thus, though Progressives tend to treat the MMR-autism connection and vaccine opposition in general as cautionary tales about the dangers of the sort of anti-scientific outlook that leads to the spread of democratically harmful misinformation, I argue the dangers of adopting the insufficiently critical attitude advocated by Progressives toward the “scientific community” are at least as important lessons to draw from those phenomena.

A similar danger is invited by Progressives' approach to the legacy news media. This is exemplified in the views of Goldman relayed in Chapter Five. As I showed in that chapter and reiterated above, Goldman defends the legacy news media on the grounds that it employs fact-checkers, limits reliance on anonymous sources, and requires multiple sources to be consulted in

forming the stories it publishes, yet offers no empirical evidence for any of these claims, instead taking them for granted. It is equally telling that Goldman does not mention the existence of countervailing incentives pushing against the rigorous employment of the processes he mentions. Fact-checking every single substantive claim made in every single article would be highly time-consuming and require large expenditures by news companies. Consulting a greater *number* and broader *diversity* of sources likewise requires reporters to do more legwork, taking up more of their time and consequently requiring the news company to pay more money. And placing strict limitations on the use of anonymous sources would require news companies to pass up potentially juicy scoops – precisely the sorts of stories most likely to capture the attention of readers who would otherwise be indifferent. Not only, then, do Goldman's claims not clearly stand up to empirical scrutiny; they also do not clearly stand up to the test of critical thinking. These oversights do not demonstrate the intellectual deficiency of Alvin Goldman, who is a thoughtful and insightful epistemologist. Instead, they demonstrate the way Progressives' institutionalism disincentivizes them from subjecting their favored institutions to the full force of their powers of critical scrutiny, causing them to content themselves with the performance of those institutions even in the absence of empirical evidence and in the face of common-sense considerations that might plausibly lead them to expect the contrary.

The key watchword here is *oversight*. The Progressive outlook does not emphasize the value of *oversight* of its favored institutions by the public. If anything, it actively opposes it. Its deeply pessimistic view of the judgment of lay citizens inclines Progressives to view lay citizens as lacking the cognitive capacity to make reliable judgments of the performance of scientists and professional journalists. This is why Progressives so often treat lay citizens' disagreement with the claims of the institutions of science or the legacy news media as evidence of a “conspiratorial,” a.k.a. *irrational*, outlook, demonizing them and *actively deterring* them from engaging in oversight and critique.⁵²⁵ So embedded is this view of lay citizens as unfit judges of the performance of certain epistemic authorities that it deters Progressives from seeing the potential value of lay citizens performing oversight even when they themselves cite examples of such oversight's being effective. This was precisely the case for Nichols, who, as I described early in this chapter, tells the story of an eighth-grader used Google to correct historians' claims

525 E.g. Lewandowsky et al (2017), Dahlgren (2018)

that “No Irish Need Apply” yet does not mention the possibility that this might demonstrate at least *some* value in lay citizens subjecting the claims of experts to scrutiny in the light of their own judgment. Epistemic pluralism resists the Progressive tendency to undervalue oversight by encouraging citizens to keep *chronic* fallibility at the forefront of their minds. Since, like Madison and Jefferson, it views fallibility as a *chronic* human condition, both in epistemic and moral terms, citizens who embrace epistemic pluralism will find it necessary to keep oversight over *all* knowledge-producing institutions. And because epistemic pluralism views each perspective as having the potential to yield meaningful insight about the world we share in common, it is less inclined to write off criticism of those institutions by lay citizens as instances of “conspiratorial” or irrational thinking.

One consequence of adopting epistemic pluralists' more charitable view of the judgmental capacities of lay citizens is that epistemic pluralists are much more likely than Progressives to interpret the much-discussed decline in public trust toward a wide range of institutions as a largely *accurate estimation* of the *actual performance* of these institutions over time. In explaining this trend, Progressives in the United States usually emphasize the anti-media and increasingly libertarian rhetoric of the political right, blaming it for – in its view, largely unjustly – undermining public trust in these institutions.⁵²⁶ While epistemic pluralists need not fully reject this claim, they would be suspicious of the idea that so many members of the public would be so consistently hoodwinked by the mere rhetoric of conservative politicians for such a long time. Instead of writing off the trend as evidence of a growing “epistemic crisis,” epistemic pluralists would be open to the idea that these trends might be indicative of *actual corruption and complacency* in the institutions in which the public consistently and increasingly expresses mistrust. This line of thinking might lead them to critically examine the performance of those institutions and the incentives placed before them. It might, for example, inspire the epistemic pluralist to look into whether the legacy news media *actually practiced* the procedures for which it is credited by Goldman, suspecting that it might find evidence that such practices were not as common as Progressives like Goldman presume.⁵²⁷ It might also lead them to treat the failure of large-scale efforts at replication in fields like psychology and cancer research, such as those

⁵²⁶ E.g. Ladd (2012), Dahlgren (2018)

⁵²⁷ I am inclined to agree with Davies (2009) and Coady (2011) that the legacy news media regularly, and increasingly, does not, in fact, employ these procedures with any great regularity.

those reported by Nichols, as sobering evidence that there might be serious problems with the incentive structures that guide the behaviors and publication processes of those institutions. That is almost certainly a more appropriate response to such findings than Nichols' own, which is to muddy the waters by insisting that such a failure need not indicate large-scale fraud, cite the arguments of researchers belonging to the same community as that being critiqued that the replication studies were “completely unfair – even irresponsible,” and offering a cursory list of “other problems” with the institutions implicated and refusing to go into any greater depth on the matter.⁵²⁸

Rather than settling for such non-explanations, citizens embracing the epistemic pluralist approach would be more inclined to respond to these findings by asking questions like:

- “How often are studies in this discipline *actually* replicated by other, independent researchers?”
- “How anonymous *is* the peer-review process in this field, usually?”
- “Are the incentives held out to reviewers likely to motivate them to do an honest, thorough and public-minded job?”
- “What are the incentives held out to researchers in this field? Are they likely to encourage them to do careful, thorough work, or are they likely to encourage them to produce results quickly and in high volume, tempting them to cut corners or even commit fraud?”
- “Are the gatekeepers in this field biased in favor of any particular point of view, or favorable only toward certain kinds of results (such as positive rather than negative findings)?”

It is very far from obvious that the public stands to lose much if increasing numbers of citizens start asking these types of questions. It is even less obvious that the institutions favored by Progressives would be able to answer them in a way that citizens ought to find satisfactory.

Another way of saying this is that it is far from obvious that the much-bemoaned decline of public trust in institutions over the last half-century is evidence of a growing epistemic “malaise” or the demagoguery of opportunistic conservative politicians. It may, on the contrary,

⁵²⁸ Nichols (2017) 184-5

be in large part a reflection of the good judgment of the public. One reason the epistemic pluralist might suspect that this is the most plausible interpretation of the public's declining trust in institutions is related to the very trend Progressives most commonly blame for *unjustly* undermining the public's trust in their favored institutions. This is, of course, the rapid expansion of information made available to lay citizens by the advance of communications technologies, beginning with cable in the 1980s and accelerated greatly by the internet and social media in the 21st century. While both Progressives and epistemic pluralists see these technological developments as likely contributors to that decline, their interpretation of that trend differs profoundly in ways that encapsulate the fundamental differences that lie between these two outlooks. Where Progressives tend to interpret declining trust in institutions as a case of the public being hoodwinked by silver-tongued politicians and other influencers given free play by an increasingly *laissez-faire* informational environment no longer watched over by its traditional “gatekeepers,” epistemic pluralists are more likely to suspect that the entry of so many new voices into the informational environment to *improve* the public's view of the world over time by giving it a broader range of points-of-view through which to view that world. No longer would citizens be limited to the information made available by a relatively small number of “gatekeepers,” most of whom held positions of relative power and privilege. Instead, now citizens would also have access to numerous perspectives beyond these power-holders. The epistemic pluralist would, therefore, suspect that the coincidence of decreased barriers to communication with declining trust in traditional institutions is in large part the result of the public's sudden ability to see the world from a wider variety of perspectives, including especially *perspectives not controlled, mediated or easily manipulated by those in power*. It would, therefore, at least consider the possibility that the public's declining trust in the legacy news media is at least partly a product of its sudden ability to compare the world as viewed through something *other than* the legacy news media against the world presented to it *by* the legacy news media, perceiving the limitations of its traditional source of information in a way it previously could not. Likewise, it would at least consider whether the public's growing mistrust of the legacy news media and other institutions was the result of the public's being able to *view* those institutions from *different angles*, allowing it to assess those institutions in a more holistic way, rather than being forced to assess them through the lens afforded them by the legacy news media

and the scant handful of other sources of information about those institutions to which it may have had access.

I believe the epistemic pluralist suspicions listed here capture much of the reason for declines in institutional trust in the internet era. Twenty-first century communications technologies have enabled citizens to see the world from different perspectives, and the world they see through them is meaningfully different from the one citizens were presented by the channels available for the bulk of the 20th century. In the world citizens see now, police brutality against Blacks is more visible. In the world citizens see now, sexual “deviance” is much more common and takes many more forms than the mainstream majority once thought. In the world citizens see now, many more people believe “radical” things than those with conventional views once believed. In the world citizens see now, mistakes and corruption among authorities are frequently brought to light and the narratives peddled by institutional authorities are regularly subjected to heated public contestation in a way that was far less likely in eras in which channels of communication were largely controlled by those same authorities or people similar to them in worldview and interests. This is but the tiniest sample of the ways in which the world viewed by citizens in the internet era poses a challenge to the world presented to citizens in the pre-internet era. The scope of that challenge is staggering and unsettling, particularly to those who put much stock in the image of the world they constructed in the pre-internet era. Its staggering and unsettling upheaval of once-accepted reality likely contributes to the increasingly divisive and fractious political climate, as suddenly citizens are forced to encounter ways of life and points of view they previously thought non-existent or marginal. To say that it partly explains the increasingly conflict-ridden discursive climate characteristic of our era, however, is a far cry from saying that the view of the world cobbled together by citizens through the internet and other 21st century communications technologies is *less accurate* than the one formed by them and their predecessors in the pre-internet era. Should we really suppose citizens' image of the world in the information-saturated internet era is inferior to that characteristic of citizens in the era prior?

I do not believe we should. Like Jefferson, I believe history clearly shows that freedoms of speech and press – the kind of “freedom of the press” meant by Jefferson – go hand-in-hand with a steady improvement in the beliefs of the public and reductions in the scope and influence of corruption. Also like Jefferson, I believe the most likely explanation for this correlation is that

the judgment of the lay public is *good enough* that, in general and over time, they are able to reliably differentiate more from less reliable claims *if allowed access to both*. If this is the case, we should expect citizens given access to a *greater quantity* and *wider variety* of information to form a correspondingly *more accurate* view of the world over time. I see no reason to suppose this is not true of 21st century democratic citizens foreboding rhetoric of Progressives who portray society as currently descending into a post-truth “malaise” notwithstanding.

Just as I agree with Jefferson's modest optimism about lay citizens, I agree with both Madison and Jefferson that all human beings individually as well as all institutions formed and administered by them are chronically fallible, both morally and epistemically. Given these facts – the *adequacy* of the public's judgment, in general and over time, and the *chronic fallibility* of all human beings and institutions created by them – the best way for us to protect our collective ability to improve our beliefs and uncover and combat illegitimate authority is for us to not only maintain the ability for multiple perspectives to proliferate in the informational environment but additionally to encourage citizens to *assertively exercise doxastic self-rule*, subjecting the claims they encounter to critical scrutiny and determining which claims to believe on the basis of their *own judgment*. This means resisting the urge to reflexively accuse citizens who refuse to trust the knowledge-disseminating institutions we personally favor of epistemic vice. It also means declining to too readily treat the adoption of views we personally find outlandish or absurd by certain citizens and groups of citizens as a “threat to democracy” or signs of a descent into a collective epistemic “malaise.”

The Public and its Surprisingly Limited Problems

In the epigraph to this chapter, John Dewey expresses optimism about the judgment of lay citizens. Dewey is often caricatured as a sort of hopeless romantic. In one respect, that reputation may be deserved. As Eric MacGilvray has argued, Dewey subscribed to a teleological view of democracy and a belief in a “genuinely common good” that can be fairly criticized as excessively romantic.⁵²⁹ But the arguments I have made in this dissertation imply that his “faith in the capacity of the intelligence of the common man” ought not be criticized by any proponent of liberal democracy on similar grounds. On the contrary, I believe Dewey is exactly correct

⁵²⁹ MacGilvray (2004), 148

when he claims that such “faith is so deeply embedded in the methods which are intrinsic to democracy that when a professed democrat denies the faith he convicts himself of treachery to his profession.”

Or, rather, *almost* exactly. Were it up to me, I would make only two amendments. The first would be to substitute the term “liberal democracy” for what Dewey calls simply “democracy,” since I have argued that a pessimist about the judgmental capacity of lay citizens could support either liberalism *or* democracy but not both. Given his consistent emphasis on the need to encourage free enquiry and freedom of communication, I believe Dewey would be friendly to this amendment. The second would be to replace the word “faith” with “belief.” The term “faith” often implies a strong conviction in a belief that cannot necessarily be supported by clear evidence, such as the existence of an eternal and omnipotent but immaterial God. I do not think the sort of modest optimism about lay citizens' judgment required to persuasively defend liberal democracy needs to be predicated on this sort of faith. On the contrary, I believe the history of contemporary democracy before and after Dewey offers plenty of support for the idea that lay citizens can be trusted, in general and over time, to embrace better ideas and resist worse ones, so long as they are allowed to be exposed to both. The extent to which this process effectively improves the collective well-being of democracy, however, depends on more than just the extent to which citizens are *allowed* to access alternative points of view. It also depends on the extent to which they are encouraged to actively *use* their powers of judgment to evaluate the claims they encounter in the informational environment instead of simply accepting or rejecting them according to whether they disseminate from or are endorsed by some pre-defined set of individuals or institutions.

Too often, contemporary political commentators view the judgment of lay citizens as a liability to democracy. What I have tried to argue in this dissertation is just the opposite. I believe the good (enough) judgment of lay citizens is one of the foundational pillars that explains the historical success of liberal democracy. It should be treated as one of its most precious assets, not one of its greatest liabilities. The communications technologies that have emerged in the 21st century offer us the opportunity to make use of that precious asset by doubling down on the “[belief] in the capacity of the intelligence of the common man” Dewey correctly identified as integral to any sound defense of [liberal] democracy. They also offer us the opportunity to

squander that potential by discouraging citizens from assertively exercising their powers of judgment or by narrowing the field of information allowed to filter down to them, thereby depriving them of “the means of judging.”⁵³⁰

It should be obvious which path I favor.

530 Mill (2007 [1859]), 78

References

- “20 Vaccine Facts You Need to Know to Make an Informed Decision.” *A Voice for Choice*. <http://avoiceforchoice.org/issues/pharmaceuticals-and-vaccines/20-vaccine-facts/> (Accessed April 2nd, 2021).
- Achen, Christopher H., and Larry M. Bartels. 2016. *Democracy for Realists : Why Elections Do Not Produce Responsive Government*. Princeton: Princeton University Press.
- Alary, Bryan. 2017. “In the Era of Fake News and Untruths, Journalism Matters.” *Folio*.
- Allcott, Hunt, and Matthew Gentzkow. 2017. “Social Media and Fake News in the 2016 Election.” *The Journal of Economic Perspectives* 31(2): 211–35.
- Alterman, Eric. 2012. *The End of Newspapers and the Decline of Democracy*. Center for American Progress.
- Ames, Gail J., and Frank B. Murray. 1982. “When Two Wrongs Make a Right: Promoting Cognitive Change by Social Conflict.” *Developmental Psychology* 18(6): 894–97.
- Anderson, Elizabeth. 2006. “The Epistemology of Democracy.” *Episteme: A Journal of Social Epistemology* 3(1): 8–22.
- Andrews, Katherine. 2017. “The Dark History of Forced Sterilization of Latina Women.” *Panoramas*. <https://www.panoramas.pitt.edu/health-and-society/dark-history-forced-sterilization-latina-women> (Accessed April 2nd, 2021).
- Apple Boss Backs Britain*. 2017. Good Morning Britain. <https://twitter.com/GMB/status/829992194366902272>.
- Arceneaux, Kevin, and Ryan J. Vander Wielen. 2017. *Taming Intuition : How Reflection Minimizes Partisan Reasoning and Promotes Democratic Accountability*. Cambridge: Cambridge University Press.
- Ari Ephraim. 2021. “Orthodox Community Faces Down Misinformation In Vaccine Education Efforts.” *Spectrum News*. <https://www.nyl.com/nyc/all-boroughs/news/2021/02/11/orthodox-community-faces-down-misinformation-in-vaccine-education-efforts> (April 2, 2021).
- Aristotle. 1962. *Nichomachean Ethics*. Indianapolis, IN: Bobbs-Merrill.

- Arrow, Kenneth. 1950. "A Difficulty in the Concept of Social Welfare." *Journal of Political Economy* 58: 328–46.
- "Art of the Lie." 2016. *The Economist*. <https://www.economist.com/leaders/2016/09/10/art-of-the-lie> (Accessed April 1st, 2021).
- Audi, Robert. 1997. "The Place of Testimony in the Fabric of Knowledge and Justification." *American Philosophical Quarterly* 34(4): 405–22.
- Bago, B., D. G. Rand, and G. Pennycook. 2020. "Fake News, Fast and Slow: Deliberation Reduces Belief in False (but Not True) News Headlines." *Journal of experimental psychology. General* 149(8): 1608–13.
- Barnard, Anne, and Michael R. Gordon. 2017. "Worst Chemical Attack in Years in Syria; U.S. Blames Assad." *The New York Times*. <https://www.nytimes.com/2017/04/04/world/middleeast/syria-gas-attack.html> (Accessed April 2, 2021).
- Barthel, Michael. 2017. "Circulation, Revenue Fall for US Newspapers Overall despite Gains for Some." *Pew Research Center*. <https://www.pewresearch.org/fact-tank/2017/06/01/circulation-and-revenue-fall-for-newspaper-industry/> (Accessed April 1st, 2021).
- Bauer, A. J., and Anthony Nadler. 2018. "Conservatives Trust Conservative Media. Here's Why." *Columbia Journalism Review*. https://www.cjr.org/tow_center/conservatives-trust-conservative-media-heres-why.php (Accessed April 2nd, 2021).
- Bean, Sandra J. 2011. "Emerging and Continuing Trends in Vaccine Opposition Website Content." *Vaccine* 29(10): 1874–80.
- Belluz, Julia. 2019. "New York's Orthodox Jewish Community Is Battling Measles Outbreaks. Vaccine Deniers Are to Blame." *Vox*. <https://www.vox.com/science-and-health/2018/11/9/18068036/measles-new-york-orthodox-jewish-community-vaccines> (Accessed April 1st, 2021).
- Bennett, W. Lance. 2016. *News : The Politics of Illusion*. 10th ed. Chicago: University of Chicago Press.
- Berelson, Bernard, Paul F. Lazarsfeld, and William N. McPhee. 1954. *Voting; a Study of Opinion Formation in a Presidential Campaign*. Chicago: University of Chicago Press.
- Bernays, Edward. 2005. *Propaganda*. Brooklyn: IG Publishing.
- Bi, Stephanie, and Tobin Klusty. 2015. "Forced Sterilizations of HIV-Positive Women: A Global Ethics and Policy Failure." *American Medical Association Journal of Ethics* 17(10): 952–57.

- Birkinshaw, Julian. 2017. "The Post-Truth World - Why Have We Had Enough Of Experts?" *Forbes*. <https://www.forbes.com/sites/lbsbusinessstrategyreview/2017/05/22/the-post-truth-world-why-have-we-had-enough-of-experts/> (Accessed April 1st, 2021).
- Blume, S. 2006. "Anti-Vaccination Movements and Their Interpretations." *Social Science & Medicine* 62(3): 628–42.
- Bond, C. F. Jr., and B. M. DePaulo. 2006. "Accuracy of Deception Judgments." *Personality and social psychology review* 10(3): 214–34.
- Boorstin, Daniel J. 1987. *The Image: A Guide to Pseudo-Events in America*. 25th Anniversary. Atheneum.
- Bord, Richard J., Ann Fisher, and Robert E. O'Connor. 1998. "Public Perceptions of Global Warming: United States and International Perspectives." *Climate Research* 11(1): 75–84.
- Boulware, L. Ebony et al. 2003. "Race and Trust in the Health Care System." *Public Health Reports* 118(4): 358–65.
- Brandon, D. T., L. A. Isaac, and T. A. LaVeist. 2005. "The Legacy of Tuskegee and Trust in Medical Care: Is Tuskegee Responsible for Race Differences in Mistrust of Medical Care?" *Journal of the National Medical Association* 97(7): 951–56.
- Brenan, Megan. 2020. *Americans Remain Distrustful of Mass Media*. Gallup. <https://news.gallup.com/poll/321116/americans-remain-distrustful-mass-media.aspx> (Accessed April 2, 2021).
- Brennan, Jason. 2016. *Against Democracy*. Princeton University Press.
- Buhler, Rich. 2018. "About." *Truth or Fiction?* <https://www.truthorfiction.com/about/> (Accessed April 1st, 2021).
- Cahill S. et al. 2017. "Stigma, Medical Mistrust, and Perceived Racism May Affect PrEP Awareness and Uptake in Black Compared to White Gay and Bisexual Men in Jackson, Mississippi and Boston, Massachusetts." *AIDS Care - Psychological and Socio-Medical Aspects of AIDS/HIV* 29(11): 1351–58.
- Calcutt, Andrew. 2016. "The Surprising Origins of 'Post-Truth' - and How It Was Spawned by the Left." *The Conversation*. <https://theconversation.com/the-surprising-origins-of-post-truth-and-how-it-was-spawned-by-the-liberal-left-68929>.
- California Democratic Party v. Jones*. 2000. 530 (U.S. Supreme Court).
- Cannon, Carl M. 2001. "The Real COMPUTER VIRUS." *American Journalism Review* 23(3).

- Caplan, Bryan Douglas 1971-. 2008. *The Myth of the Rational Voter : Why Democracies Choose Bad Policies*. Princeton, NJ: Princeton University Press.
- Carson, James, and Michael Cogley. 2021. "Fake News: What Exactly Is It – and How Can You Spot It?" *The Telegraph*. <https://www.telegraph.co.uk/technology/0/fake-news-exactly-donald-trump-rise/> (Accessed April 2nd, 2021).
- Chakravartty, Anjan. 2017. "Scientific Realism." In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. <https://plato.stanford.edu/archives/sum2017/entries/scientific-realism/> (Accessed April 2nd, 2021).
- Chambers, Simone. 2010. "Theories of Political Justification." *Philosophy Compass* 5(11): 893–903.
- Chokshi, Niraj, and Julia Jacobs. 2019. "24 Charged in \$1.2 Billion Medicare Scheme, U.S. Says." *The New York Times*.
- Coady, C. A. J. 1992. *Testimony : A Philosophical Study*. Oxford: Clarendon Press.
- Coady, David 1965-. 2012. *What to Believe Now : Applying Epistemology to Contemporary Issues*. Chichester, West Sussex: Wiley-Blackwell.
- Collins, Henry, and Trevor Pinch. 2007. *The Golem: What You Should Know About Science*. 2nd ed. Cambridge University Press.
- Collins, Megan Eileen, and Thomas A. Loughran. 2017. "Rational Choice Theory, Heuristics, And Biases." In *The Oxford Handbook of Offender Decision Making*, Oxford University Press.
- Conant, James. 2006. "The Dialectic of Perspectivism, II." *Sats - Nordic Journal of Philosophy* 7(1).
- "Continental Congress to the Inhabitants of Quebec." 1774.
- Converse, Philip E. 1964. "The Nature of Belief Systems in Mass Publics." *Critical Review* 18(1–3): 1–74.
- Cook, Timothy E. 1998. *Governing with the News : The News Media as a Political Institution*. Chicago: University of Chicago Press.
- Crouse, Timothy. 1973. *The Boys on the Bus*. New York: Random House.
- Dahl, Robert. 2006. *Preface to Democratic Theory: Expanded Edition*. University of Chicago Press.

- Dahl, Robert A. 2005. "James Madison: Republican or Democrat?" *Perspectives on Politics* 3(3): 439–48.
- Dahlgren, Peter. 2018. "Media, Knowledge and Trust: The Deepening Epistemic Crisis of Democracy." *Javnost - The Public* 25(1–2): 20–27.
- Daston, Lorraine 1951-, and Peter 1955- Galison. 2007. *Objectivity*. New York: Zone Books.
- Davies, Nick 1953-. 2008. *Flat Earth News : An Award-Winning Reporter Exposes Falsehood, Distortion and Propaganda in the Global Media*. London: Chatto & Windus.
- Davies, William. 2016. "The Age of Post-Truth Politics." *The New York Times*.
- Deb, Anamitra, Stacy Donohue, and Tom Glaisyer. 2017. *Is Social Media a Threat to Democracy?* The Omidyar Group.
- Deer, Brian. 2007. "What Makes an Expert?" *British Medical Journal* 334(7595): 666–67.
- . 2011a. "How the Case Against the MMR Vaccine Was Fixed." *British Medical Journal* 342(7788): 77–82.
- . 2011b. "How the Vaccine Crisis Was Meant to Make Money." *British Medical Journal* 342(7789): 136–42.
- Dekker, Paul., and Eric M. Uslaner. 2001. *Social Capital and Participation in Everyday Life*. London: Routledge.
- Delli Carpini, Michael X. 1953-, and Scott. Keeter. 1996. *What Americans Know about Politics and Why It Matters*. New Haven: Yale University Press.
- Derakhshan, Hossein. 2019. "If News Is Dying, Who Will Safeguard Democracy?" *The Guardian*.
- Dewey, John. 1976. "Creative Democracy: The Task Before Us." In *John Dewey: The Later Works, 1925-1953*, Southern Illinois University Press, 224–30.
- Digital, Culture, Media and Sport Committee. 2018. *Disinformation and "Fake News": Interim Report*. House of Commons.
- Dill, William A. 1928. "Growth of Newspapers in the United States." University of Oregon. Masters Thesis.
- Douven, Igor. 2019. "Optimizing Group Learning: An Evolutionary Computing Approach." *Artificial Intelligence* 275: 235–51.
- Downs, Anthony. 1965. *An Economic Theory of Democracy*. New York: Harper & Row.

- Dubois, Elizabeth, and Grant Blank. 2018. "The Echo Chamber Is Overstated: The Moderating Effect of Political Interest and Diverse Media." *Information, Communication & Society* 21(5): 729–45.
- Dugan, Andrew. 2014. "Americans' Confidence in News Media Remains Low." *Gallup.com*. <https://news.gallup.com/poll/171740/americans-confidence-news-media-remains-low.aspx> (Accessed April 2nd, 2021).
- "Editorial: The Dangers of Fake News." 2016. *Tampa Bay Times*. <https://www.tampabay.com/opinion/editorials/editorial-the-dangers-of-fake-news/2304118/> (Accessed April 2nd, 2021).
- El-Bermawy, Mostafa M. 2016. "Your Echo Chamber Is Destroying Democracy | WIRED." *Wired*. <https://www.wired.com/2016/11/filter-bubble-destroying-democracy/> (Accessed April 2nd, 2021).
- Esposito S et al. 2019. "Pertussis Prevention: Reasons for Resurgence, and Differences in the Current Acellular Pertussis Vaccines." *Frontiers in immunology* 10: 1344.
- Estlund, David M. 2009. *Democratic Authority : A Philosophical Framework*. Princeton: Princeton University Press.
- Even. 2020. "On Why The Social Justice Warrior Ideology Is Fundamentally Medieval." *Medium*. <https://medium.com/@even.aesphasian/on-why-the-social-justice-warrior-ideology-is-fundamentally-medieval-fd6a09a2dcc8> (Accessed April 2nd, 2021).
- "Facebook and the Digital Virus Called Fake News." 2016. *The New York Times*. <https://www.nytimes.com/2016/11/20/opinion/sunday/facebook-and-the-digital-virus-called-fake-news.html> (Accessed April 2nd, 2021).
- Farkas, Johan, and Jannick Schou. 2019. *19 Post-Truth, Fake News and Democracy : Mapping the Politics of Falsehood*. Abingdon, Oxon: Routledge.
- Fishkin, James S. 2009. *When the People Speak : Deliberative Democracy and Public Consultation*. Oxford: Oxford University Press.
- Flaxman S., Goel S., and Rao J.M. 2016. "Filter Bubbles, Echo Chambers, and Online News Consumption." *Public Opinion Quarterly* 80(Specialissue1): 298–320.
- Fletcher, Richard, Alessio Cornia, Lucas Graves, and Rasmus Kleis Nielsen. 2018. *Measuring the Reach of "Fake News" and Disinformation in Europe*. Reuters Institution & the University of Oxford.
- Fricker, Elizabeth. 1995. "Critical Notice." *Mind* 104(414): 393–411.
- Fuller, Steve. 2018. *Post-Truth: Knowledge as a Power Game*. London: Anthem Press.

- Galison, Peter. 2000. "Einstein's Clocks: The Place of Time." *Critical Inquiry* 26(2): 355–89.
- Gamble, Vanessa Northington. 1993. "A Legacy of Distrust: African Americans and Medical Research." *American Journal of Preventive Medicine* 9(6): 35–38.
- Gangarosa EJ et al. 1998. "Impact of Anti-Vaccine Movements on Pertussis Control: The Untold Story." *Lancet (London, England)* 351: 356–61.
- Garrett, R. Kelly. 2017. "The 'Echo Chamber' Distraction: Disinformation Campaigns Are the Problem, Not Audience Fragmentation." *Journal of Applied Research in Memory and Cognition* 6(4): 370–76.
- Gentzkow, Matthew, and Jesse M. Shapiro. 2011. "Ideological Segregation Online and Offline." *The Quarterly Journal of Economics* 126(4): 1799–1839.
- Gerken, Heather K. 2005. "Second-Order Diversity." *Harvard Law Review* 118(4): 1099–1196.
- Globla Risks 2013: Eighth Edition*. 2013. World Economic Forum.
- Goldman, Alvin. 2008. "The Social Epistemology of Blogging." In *Information Technology and Moral Philosophy*, eds. Jeroen Van Den Hoven and John Weckert. Cambridge University Press, 111–22.
- Goldman, Alvin, and Bob Beddor. 2016. "Reliabilist Epistemology." In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2016/entries/reliabilism/> (Accessed April 2nd, 2021).
- Graber, Doris. 2003. "The Media and Democracy: Beyond Myths and Stereotypes." *Annual Review of Political Science* 6(1).
- Guess, Andrew, Jonathan Nagler, and Joshua Tucker. 2019. "Less Than You Think: Prevalence and Predictors of Fake News Dissemination on Facebook." *Science Advances* 5.
- Hannan J. 2018. "Trolling Ourselves to Death? Social Media and Post-Truth Politics." *European Journal of Communication* 33(2): 214–26.
- Hardwig, John. 1985. "Epistemic Dependence." *The Journal of Philosophy* 82(7): 335–49.
- Hayek, Friedrich A. 2011. *The Constitution of Liberty*. University of Chicago Press.
- Hazard Owen, Laura. 2020. "Republicans and Democrats Live in 'Nearly Inverse News Media Environments,' Pew Finds." *Nieman Lab*. <https://www.niemanlab.org/2020/01/republicans-and-democrats-live-in-nearly-inverse-news-media-environments-pew-finds/> (Accessed April 2nd, 2021).

- Hazlett, Thomas W., and David W. Sosa. 1997. "Was The Fairness Doctrine A 'Chilling Effect'? Evidence from the Postderegulation Radio Market." *The Journal of Legal Studies* 26(1): 279–301.
- Hersh, Seymour M. 2013. "Whose Sarin?" *London Review of Books* 35(24).
<https://www.lrb.co.uk/the-paper/v35/n24/seymour-m.-hersh/whose-sarin> (Accessed April 2nd, 2021).
- Hetherington, Marc, and Jonathan M. Ladd. 2020. "Destroying Trust in the Media, Science, and Government Has Left America Vulnerable to Disaster." *Brookings*.
<https://www.brookings.edu/blog/fixgov/2020/05/01/destroying-trust-in-the-media-science-and-government-has-left-america-vulnerable-to-disaster/> (Accessed April 2nd, 2021).
- Hill, Holly A. et al. 2018. *Vaccination Coverage Among Children Aged 19-35 Months - United States 2017*. US Department of Health and Human Services/Centers for Disease Control and Prevention.
- Hinderaker, John. 2004. "The Post Recaps Rathergate." *Power Line*.
<https://www.powerlineblog.com/archives/2004/09/007827.php> (Accessed April 2nd, 2021).
- "How Google, Facebook and Others Are Using Algorithms to Fight Coronavirus Misinformation." 2020. *Economic Times Telecom*.
<https://telecom.economictimes.indiatimes.com/news/how-google-facebook-and-others-are-using-algorithms-to-fight-coronavirus-misinformation/75214933> (Accessed April 2nd, 2021).
- Hume, David. 1955. *An Inquiry Concerning Human Understanding*. Bobbs-Merrill.
- Huss, Nick. 2020. "How Many Websites Are There Around the World?" *Siteefy*.
<https://siteefy.com/how-many-websites-are-there/> (Accessed April 1st, 2021).
- Indicators of News Media Trust*. 2018. Gallup/Knight Foundation.
- "Initiatives to Counter Fake News." 2019. *Library of Congress*.
<https://www.loc.gov/law/help/fake-news/france.php> (Accessed April 2nd, 2021).
- Jackson, Dean. 2018. *Issue Brief: How Disinformation Impacts Politics and Publics*. National Endowment for Democracy.
- Jaiswal, Jessica, and Perry N. Halkitis. 2019. "Towards a More Inclusive and Dynamic Understanding of Medical Mistrust Informed by Science." *Behavioral Medicine* 45(2): 79–85.

- James, William. 1968. "What Pragmatism Means." In *Pragmatism and Other Essays*, Washington Square Press, 22–38.
- . 1981. *The Principles of Psychology*. Cambridge, MA: Harvard University Press.
- Jamieson, Amber. 2016. "DNC Head Leaked Debate Question to Clinton, Podesta Emails Suggests." *The Guardian*. <http://www.theguardian.com/us-news/2016/oct/31/donna-brazile-hillary-clinton-debate-question-podesta-emails-cnn> (Accessed April 2nd, 2021).
- Jamieson, Kathleen Hall. 2015. "Implications of the Demise of 'Fact' in Political Discourse." *Proceedings of the American Philosophical Society* 159(1): 66–84.
- . 2018. *Cyberwar : How Russian Hackers and Trolls Helped Elect a President : What We Don't, Can't, and Do Know*. New York, NY: Oxford University Press.
- Jamieson, Kathleen Hall., and Paul. Waldman. 2003. *The Press Effect : Politicians, Journalists, and the Stories That Shape the Political World*. Oxford ; Oxford University Press.
- Jefferson, Thomas. 1805. "2nd Inaugural Address." Presented at the Washington, D.C.
- Jones, Jeff, and Saad, Lydia. 2017. "Gallup Poll Series: Environment."
- Jones, Jeffrey M. 2018. "U.S. Media Trust Continues to Recover From 2016 Low." *Gallup*. <https://news.gallup.com/poll/243665/media-trust-continues-recover-2016-low.aspx> (Accessed April 2nd, 2021).
- Journalism, "Fake News" & Disinformation: Handbook for Journalism Education and Training*. 2018. UNESCO.
- "Journalists Are Not the Enemy." 2018. *The Boston Globe*. <https://www.bostonglobe.com/2018/08/15/editorial/MRD1jQwctHr3Q3y3QkRDfK/story.html> (Accessed April 1st, 2021).
- Jurkowitz, Mark, and Amy Mitchell. 2020. *An Oasis of Bipartisanship: Republicans and Democrats Distrust Social Media Sites for Political and Election News*. Pew Research Center.
- Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. New York: Anchor Canada.
- Kanai K. 1980. "Japan's Experience in Pertussis Epidemiology and Vaccination in the Past Thirty Years." *Japanese journal of medical science & biology* 33(3): 107–43.
- Kant, Immanuel. 1991. "Idea for a Universal History with a Cosmopolitan Purpose." In *Kant: Political Writings*, Cambridge University Press, 41–53.

- Karlova, Natascha A., and Karen E. Fisher. 2012. “‘Plz RT’: A Social Diffusion Model of Misinformation and Disinformation for Understanding Human Information Behaviour.” Presented at the International Symposium on Information and Communication 2012, Tokyo.
- Katz, Elihu. 1957. “The Two-Step Flow of Communication: An Up-To-Date Report on an Hypothesis.” *The Public Opinion Quarterly* 21(1): 61–78.
- Kelly, Jamie Terence. 2012. *Framing Democracy : A Behavioral Approach to Democratic Theory*. Princeton, NJ: Princeton University Press.
- Kitcher, Philip. 1990. “The Division of Cognitive Labor.” *The Journal of Philosophy* 87(1): 5–22.
- Klawitter, Erin, and Eszter Hargittai. 2018. “Shortcuts to Well Being? Evaluating the Credibility of Online Health Information through Multiple Complementary Heuristics.” *Journal of Broadcasting & Electronic Media* 62(2): 251–68.
- Kleemans, M. et al. 2009. “Sensationalism in Television News: A Review.” In *Meaningful Media: Communication Research on the Social Construction of Reality*, ed. F. J. M. Huysmans. Tandem Felix, 226–43.
- Klein, Naomi. 2007. *The Shock Doctrine : The Rise of Disaster Capitalism*. 1st ed. New York: Metropolitan Books/Henry Holt.
- Kovach, Bill, and Tom Rosenstiel. 2011. *Blur: How to Know What’s True in the Age of Information Overload*. Bloomsbury Publishing.
- Kuhn, Thomas. 2012. *The Structure of Scientific Revolutions*. 4th ed. University of Chicago Press.
- Kuklinski, James H. et al. 2000. “Misinformation and the Currency of Democratic Citizenship.” *The Journal of Politics* 62(3): 790–816.
- Lackey, Jennifer. 2003. “A Minimal Expression of Non-Reductionism in the Epistemology of Testimony.” *Noûs* 37(4): 706–23.
- . 2011. “Testimony: Acquiring Knowledge from Others.” In *Social Epistemology: Essential Readings*, eds. Alvin Goldman and Dennis Whitcomb. Oxford University Press.
- Ladd, Jonathan M. 2012. *Why Americans Hate the Media and How It Matters*. Princeton University Press.
- Lakatos, Imre 1922-1974. 1978. *The Methodology of Scientific Research Programmes*. Cambridge: Cambridge University Press.

- Landemore, Hélène. 2012. *Democratic Reason : Politics, Collective Intelligence, and the Rule of the Many*. Princeton: Princeton University Press.
- Larson HJ et al. 2011. “Addressing the Vaccine Confidence Gap.” *Lancet (London, England)* 378(9790): 526–35.
- LaVeist, T. A. et al. 2009. “Environmental and Socio-Economic Factors as Contributors to Racial Disparities in Diabetes Prevalence.” *Journal of general internal medicine* 24(10): 1144–48.
- Lawrence, Jane. 2000. “The Indian Health Service and the Sterilization of Native American Women.” *American Indian Quarterly* 24(3): 400–419.
- Lazer, David M. J. et al. 2018. “The Science of Fake News: Addressing Fake News Requires a Multidisciplinary Effort.” *Science* 359(6380).
- Lederman, Reeva, Hanmei Fan, Stephen Smith, and Shanton Chang. 2014. “Who Can You Trust? Credibility Assessment in Online Health Forums.” *Health Policy and Technology* 3(1): 13–25.
- Levinson, Jessica Levinson Jessica. 2019. “Opinion | Politicians Shouldn’t Be Able to Lie on Facebook. Here’s How to Stop Them.” *NBC News*.
<https://www.nbcnews.com/think/opinion/facebook-has-political-fake-news-problem-can-we-fix-it-ncna1071001> (Accessed April 2nd, 2021).
- Levitsky, Steven, and Daniel Ziblatt. 2018. *How Democracies Die*. First edition. New York: Crown. <http://worldcat.org/oclc/986837776>.
- Lewandowsky, S. et al. 2012. “Misinformation and Its Correction: Continued Influence and Successful Debiasing.” *Psychological science in the public interest : a journal of the American Psychological Society* 13(3): 106–31.
- Lewandowsky, Stephan. 2019. “In Whose Hands the Future? Confronting Climate Science Denial and Conspiracy Theories.” Presented at the Summer Institute on Bounded Rationality, Berlin, Germany. <https://www.youtube.com/watch?v=DHdqv1lyZgg>
- Lewandowsky, Stephan, John Cook, and Ullrich K.H. Ecker. 2017. “Letting the Gorilla Emerge From the Mist: Getting Past Post-Truth.” *Journal of Applied Research in Memory and Cognition* 6(4): 418–24.
- Lewandowsky, Stephan, Ullrich K.H. Ecker, and John Cook. 2017. “Beyond Misinformation: Understanding and Coping with the ‘Post-Truth’ Era.” *Journal of Applied Research in Memory and Cognition* 6(4): 353–69.
- Liebling, A. J. 1947. *The Wayward Pressman*. 1st ed. Garden City, NY: Doubleday.

- Lippmann, Walter. 1920. *Liberty and the News*. BN Publishing.
- . 1925. *The Phantom Public*. New York, NY: Harcourt, Brace and Company.
- . 2012. *Public Opinion*. Middletown, DE: Renaissance Books.
- Locke, John. 1980. *Second Treatise of Government*. Hackett.
- Lodge, Milton, and C. S. Taber. 2013. *The Rationalizing Voter*. Cambridge University Press.
- Loewenstein, Antony. 2008. *The Blogging Revolution*. Melbourne University Press.
- Lucassen, Teun, Rienco Muilwijk, Matthijs L. Noordzij, and Jan Maarten Schraagen. 2013. “Topic Familiarity and Information Skills in Online Credibility Evaluation.” *Journal of the American Society for Information Science and Technology* 64(2): 254–64.
- Lynn, L. E. 2011. “Federalist No. 51: Is Liberty Guaranteed by Structures?” *Public Administration Review* 71(Supplementary Issue 1): S83–89.
- M Vav. 2019. *Shar Yaqataz Banyamyan Facebook Video*. https://www.youtube.com/watch?t=4340&v=t3EC1_gcr34&feature=youtu.be (Accessed April 2nd, 2021).
- MacCoun, Robert J. 1998. “Biases in the Interpretation and Use of Research Results.” *Annual Review of Psychology* 49(1).
- MacGilvray, Eric. 2004. *Reconstructing Public Reason*. Harvard University Press.
- Madison, James. 1787. “Federalist No. 10.”
- . 1788a. “Federalist No. 47.”
- . 1788b. “Federalist No. 51.”
- Mandeville, Bernard. 1962. *The Fable of the Bees : Or, Private Vices, Publick Benefits*. New York: Capricorn Books.
- Manjoo, Farhad. 2008. *True Enough : Learning to Live in a Post-Fact Society*. Hoboken, N.J.: Wiley.
- Mason, Winter, and Duncan J. Watts. 2012. “Collaborative Learning in Networks.” *Proceedings of the National Academy of Sciences of the United States of America* 109(3): 764–69.
- Mayo-Wilson, Conor, Kevin J. S. Zollman, and David Danks. 2011. “The Independence Thesis: When Individual and Social Epistemology Diverge*.” *Philosophy of Science* 78(4): 653–77.

- McDonough, Katie. 2013. "Israel Admits Ethiopian Jewish Immigrants Were given Birth Control Shots." *Salon*.
https://www.salon.com/2013/01/28/israel_admits_ethiopian_jewish_immigrants_were_given_birth_control_shots/ (Accessed April 2nd, 2021).
- McGrew, Sarah, Teresa Ortega, Joel Breakstone, and Sam Wineburg. 2017. "The Challenge That's Bigger than Fake News: Civic Reasoning in a Social Media Environment." *American Educator* Fall: 4–39.
- McGrew, Timothy. 2003. "Confirmation, Heuristics, and Explanatory Reasoning." *The British Journal for the Philosophy of Science* 54(4): 553–67.
- "Measles 2019: Washington State Department of Health." 2019. *Washington State Department of Health*.
<https://www.doh.wa.gov/YouandYourFamily/IllnessandDisease/Measles/Measles2019>
 (Accessed April 1st, 2021).
- Mervosh, Sarah, and Emily S. Rueb. 2019. "Fuller Picture Emerges of Viral Video of Native American Man and Catholic Students." *The New York Times*.
<https://www.nytimes.com/2019/01/20/us/nathan-phillips-covington.html> (Accessed April 2nd, 2021).
- Metzger, Miriam J., Andrew J. Flanagin, and Ryan B. Medders. 2010. "Social and Heuristic Approaches to Credibility Evaluation Online." *Journal of Communication* 60(3): 413–39.
- Mill, John Stuart. 2007. *On Liberty*. Pearson Longman.
- Miller, Alexander. 2019. "Realism." In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Stanford University. <https://plato.stanford.edu/archives/win2019/entries/realism/>
 (Accessed April 2nd, 2021).
- Mills, Curt. 2017. "Syria: Narrative That Assad Behind Chemical Attack Questioned | World Report | US News." *U.S. News and World Report*.
<https://www.usnews.com/news/world/articles/2017-04-06/syria-narrative-that-assad-behind-chemical-attack-questioned> (Accessed April 2nd, 2021).
- Minozzi, William. 2011. "A Jamming Theory of Politics." *Journal of Politics* 73(2).
- Minozzi, William, and Michael Neblo. 2019. "Thrasymachus's Blush: Between the Science of Reason & the Politics of Emotion."
- Mitchell, Amy. 2014. *8 Key Takeaways About Social Media and News*. Pew Research Center.
- Miztal, Barbara. 2001. "Trust and Cooperation: The Democratic Public Sphere." *Journal of Sociology* 37(4): 371–86.

- Moeller, Susan D. 2004. "Follow the Leader: In Reporting on Weapons of Mass Destruction, the Media Too Often Take Their Cues from the President.(First Person)." *American Journalism Review* 28(2): 16–31.
- Montesquieu, Baron de. 1952. 38 *The Spirit of Laws*. ed. William Benton. Encyclopaedia Britannica, Inc.
- Muldoon, Ryan. 2016. *Social Contract Theory for a Diverse World: Beyond Tolerance*. New York, NY: Routledge-Taylor & Francis.
- Nagel, Thomas. 1987. "Moral Conflict and Political Legitimacy." *Philosophy & Public Affairs* 16(3): 215–40.
- Newman, Nic, and Richard Fletcher. 2017. "Bias, Bullshit and Lies: Audience Perspectives on Low Trust in the Media." *SSRN Electronic Journal*.
- Newton, Kenneth. 2001. "Trust, Social Capital, Civil Society, and Democracy." *International Political Science Review* 22(2): 201–14.
- Nichols, Thomas M. 2017. *The Death of Expertise : The Campaign against Established Knowledge and Why It Matters*. New York, NY: Oxford University Press.
- Nisbet, Matthew C., and Teresa Myers. 2007. "Twenty Years of Public Opinion about Global Warming." *Public Opinion Quarterly* 71(3): 444–70.
- Nixon v. Shrink Missouri Government*. 2000. 528 (U.S. Supreme Court).
- Nurse, Matthew S., and Will J. Grant. 2020. "I'll See It When I Believe It: Motivated Numeracy in Perceptions of Climate Change Risk." *Environmental Communication: A Journal of Nature and Culture* 14(2): 184–201.
- Nyhan, Brendan. 2018. "Fake News Is Troubling, But Its Effect Is Limited." *New York Times* 167(57873).
- "Nytimes.Com Traffic, Ranking & Marketing Analytics." 2021. *SimilarWeb*. <http://similarweb.com/website/nytimes.com/> (Accessed April 1st, 2021).
- O'Donnell, Norah. 2020. "The Good Doctor." *InStyle*. <https://www.instyle.com/news/dr-fauci-says-with-all-due-modesty-i-think-im-pretty-effective> (Accessed April 2nd, 2021).
- Ortiz-Ospina, Esteban. 2019. "The Rise of Social Media." *Our World in Data*. <https://ourworldindata.org/rise-of-social-media> (Accessed April 1st, 2021).
- Ovadya, Aviv. 2018. "What's Worse than Fake News? The Distortion of Reality Itself." *New Perspectives Quarterly* 35(2): 43–45.

- Page, Scott E. 2007. *The Difference : How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton: Princeton University Press.
- Paris, Francesca. 2019. "Video Of Kentucky Students Mocking Native American Man Draws Outcry : NPR." *NPR*. <https://www.npr.org/2019/01/20/686988268/video-of-kentucky-students-mocking-native-american-man-draws-outcry> (Accessed April 2nd, 2021).
- Pariser, Eli. 2011. *The Filter Bubble : What the Internet Is Hiding from You*. New York: Penguin Press.
- Patel, Manisha, Howard Zucker, Oxiris Barbot, and Amanda Cohn. 2020. "Measles – Maintaining Disease Elimination and Enhancing Vaccine Confidence." Presented at the Centers for Disease Control and Prevention. <https://www.cdc.gov/grand-rounds/pp/2020/20200218-measles-elimination-vaccine.html> (Accessed April 2nd, 2021).
- Pear, Robert. 2005. "Bush Team Is Criticized for 'Covert Propaganda.'" *The New York Times*. <https://www.nytimes.com/2005/10/02/world/americas/bush-team-is-criticized-for-covert-propaganda.html> (Accessed April 2nd, 2021).
- Peirce, Charles Sanders. 2006. "Some Consequences of Four Incapacities." In *Pragmatism: Old & New*, ed. Susan Haack. Prometheus Books, 69–105.
- Pennycook, Gordon, and David G. Rand. 2019. "Lazy, Not Biased: Susceptibility to Partisan Fake News Is Better Explained by Lack of Reasoning than by Motivated Reasoning." *Cognition* 188: 39–50.
- Persily, Nathaniel. 2017. "Can Democracy Survive the Internet?" *Journal of Democracy* 28(2): 63–76.
- Pilkington, Ed. 2018. "Trump v CNN: Lawsuit Becomes Test Case on Press Freedom." *The Guardian*. <http://www.theguardian.com/media/2018/nov/14/trump-cnn-jim-acosta-lawsuit-whats-at-stake-press-freedom> (Accessed April 2nd, 2021).
- Plato. 1968. *The Republic of Plato*. 2nd ed. Basic Books.
- Popkin, Samuel L. 1991. *The Reasoning Voter : Communication and Persuasion in Presidential Campaigns*. Chicago: University of Chicago Press.
- Popper, Karl R. (Karl Raimund), 1902-1994. 1971. *The Open Society and Its Enemies*. 5th ed. Princeton, NJ: Princeton University Press.
- Posner, Richard A. 2003. *Law, Pragmatism, and Democracy*. Cambridge, MA: Harvard University Press.
- . 2005. "Bad News." *New York Times Book Review*: 1, 8–11.

- “‘Post-Truth’ Declared Word of the Year by Oxford Dictionaries.” 2016. *BBC News*.
<https://www.bbc.com/news/uk-37995600> (Accessed April 1st, 2021).
- Price, Daniel. 2006. “Fake TV News: Video News Releases.” *PR Watch*.
<https://www.prwatch.org/fakenews/findings/vnrs> (Accessed April 2nd, 2021).
- Prior, Markus 1974-. 2007. *Post-Broadcast Democracy : How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*. New York: Cambridge University Press.
- Przeworski, Adam. 1999. “Minimalist Conception of Democracy: A Defense.” In *Democracy’s Value*, Contemporary Political Theory, Cambridge University Press, 12–17.
- Putnam, Robert D. 2000. *Bowling Alone : The Collapse and Revival of American Community*. New York: Simon & Schuster.
- Rafferty, rew. 2017. “First Amendment Advocates Push Back on Trump’s Licensing Threat.” *NBC News*. <https://www.nbcnews.com/politics/politics-news/first-amendment-advocates-push-back-trump-s-licensing-threat-n809941> (Accessed April 2nd, 2021).
- Randall, Eric. 2011. “Media Non-Coverage of Occupy Wall Street Gets Lots of Media Coverage.” *The Atlantic*. <https://www.theatlantic.com/national/archive/2011/09/media-non-coverage-occupy-wall-street-gets-lots-media-coverage/337391/> (Accessed April 2nd, 2021).
- Rao, T. S. Sathyanarayana, and Chittaranjan Andrade. 2011. “The MMR Vaccine and Autism: Sensation, Refutation, Retraction, and Fraud.” *Indian Journal of Psychiatry* 53(2): 95–96.
- Rawls, John. 2005. *Political Liberalism: Expanded Edition*. New York: Columbia University Press.
- Reid, Thomas. 1983. *Thomas Reid’s Inquiry and Essays*. eds. Ronald E. Beanblossom and Keith Lehrer. Hackett.
- Resnick, Brian. 2017. “The Science behind Why Fake News Is so Hard to Wipe Out.” *Vox*.
<https://www.vox.com/science-and-health/2017/10/5/16410912/illusory-truth-fake-news-las-vegas-google-facebook> (Accessed April 2nd, 2021).
- Rorty, Richard. 1991. “Cosmopolitan without Emancipation.” In *Objectivity, Relativism, and Truth: Philosophical Papers Volume I*, Cambridge University Press, 211–22.
- Rosenstock, Sarita, Justin Bruner, and Cailin O’Connor. 2017. “In Epistemic Networks, Is Less Really More?” *Philosophy of Science* 84(2): 234–52.
- Roser, Max, and E. Ortiz-Ospina. 2016. “Global Rise of Education.” *Our World in Data*.
 (Accessed April 1st, 2021).

- Roth, Maddy. 2012. "Facts Still Dead." *American Journalism Review* Winter: 4–5.
- Rothstein, Aaron. 2015. "Vaccines and Their Critics, Then and Now." *The New Atlantis: A Journal of Technology & Society*: 3–27.
- Rubin, Olivier. 2009. "The Merits of Democracy in Famine Protection – Fact or Fallacy?" *European Journal of Development Research* 21(5).
- Said-Moorhouse, Lauren, and Sarah Tilotta. 2017. "Airstrike to US Intervention: How Attack Unfolded." *CNN*. <https://www.cnn.com/2017/04/05/middleeast/syria-airstrike-idlib-how-it-unfolded/index.html> (Accessed April 2nd, 2021).
- Sales, Ben. 2019. "I Attended an Orthodox Anti-Vaccine Rally. Here's What I Saw." *Jewish Telegraphic Agency*. <https://www.jta.org/2019/06/05/united-states/an-orthodox-rally-in-brooklyn-sees-vaccines-as-a-conspiracy> (Accessed April 2nd, 2021).
- Satariano, Adam. 2020. "Republican Distrust of News May Be Helping Election Misinformation Spread." *The New York Times*. <https://www.nytimes.com/2020/11/12/technology/republicans-distrust-news-media.html> (Accessed April 2nd, 2021).
- Schattschneider, E. E. 1960. *The Semisovereign People*. Hinsdale, IL: Dryden Press.
- Schiffrin, Anya. 2017. "Disinformation and Democracy - The Internet Transformed Protest but Did Not Improve Democracy." *Journal of International Affairs* 71(1): 117–26.
- Schudson, Michael. 1978. *Discovering the News: A Social History of American Newspapers*. New York: Basic Books.
- . 1990. "Preparing the Minds of the People: Three Hundred Years of the American Newspaper." Presented at the American Antiquarian Society Seminar Commemorating 300 Years of the American Newspaper, Gannett Foundation Freedom Center.
- . 1998. *The Good Citizen: A History of American Civic Life*. New York: Martin Kessler Books.
- Schumpeter, Joseph. 1976. *Capitalism, Socialism, and Democracy*. Harper & Row.
- Sen, Amartya Kumar. 1999. "Democracy as a Universal Value." *Journal of Democracy* 10(3): 3–17.
- Shapiro, Robert, and Yaeli Bloch-Elkon. 2008. "Do the Facts Speak for Themselves? Partisan Disagreement as a Challenge to Democratic Competence." *Critical Review* 20(1–2): 115–39.

- Shearer, Elisa, and Katerina Eva Matsa. 2018. "News Use Across Social Media Platforms 2018." *Pew Research Center's Journalism Project*. <https://www.journalism.org/2018/09/10/news-use-across-social-media-platforms-2018/> (Accessed April 2nd, 2021).
- Silverman, Craig. 2015. *Lies, Damn Lies, and Viral Content: How News Websites Spread (and Debunk) Online Rumors, Unverified Claims, and Misinformation*. Tow Center for Digital Journalism.
- . 2016. "This Analysis Shows How Viral Fake Election News Stories Outperformed Real News On Facebook." *Buzzfeed*. <https://www.buzzfeednews.com/article/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook> (Accessed April 1st, 2021).
- Sirucek, Stefan. 2010. "An Interview With Howard Zinn." *The Huffington Post*. https://www.huffpost.com/entry/an-interview-with-howard_b_440159 (Accessed April 2nd, 2021).
- Smith, Adam. 1937. *The Wealth of Nations*. The Modern Library.
- Smith, Matthew. 2019. "Britons Least Likely of 22 Nations to Trust Information on Social Media | YouGov." *YouGov*. <https://yougov.co.uk/topics/technology/articles-reports/2019/05/07/britons-least-likely-22-nations-trust-information-> (Accessed April 2nd, 2021).
- "Social Media & COVID-19: A Global Study of Digital Crisis Interaction among Gen Z and Millennials." 2021. *World Health Organization*. <https://www.who.int/news-room/feature-stories/detail/social-media-covid-19-a-global-study-of-digital-crisis-interaction-among-gen-z-and-millennials> (Accessed April 2nd, 2021).
- Somin, Ilya. 2016. *Democracy and Political Ignorance : Why Smaller Government Is Smarter*. 2nd ed. Stanford, California: Stanford Law Books, an imprint of Stanford University Press.
- Soroka, Stuart N. 2012. "The Gatekeeping Function: Distributions of Information in Media and the Real World." *The Journal of Politics* 74(2): 514–28.
- Sparrow, Bartholomew H. 1959-. 1999. *Uncertain Guardians : The News Media as a Political Institution*. Baltimore: Johns Hopkins University Press.
- Speed E and Mannion R. 2017. "The Rise of Post-Truth Populism in Pluralist Liberal Democracies: Challenges for Health Policy." *International journal of health policy and management* 6(5): 249–51.
- Spivak, Gayatri Chakravorty. 1988. "Can the Subaltern Speak?" In *Marxism and the Interpretation of Culture*, University of Illinois Press.

- Starr, Paul 1949-. 2004. *The Creation of the Media : Political Origins of Modern Communications*. New York: Basic Books.
- State of the News: 2004*. 2004. Pew Research Center.
- Stegmaier, Mary, Michael S. Lewis-Beck, and Beomseob Park. 2016. "The VP-Function: A Review." In *The Sage Handbook of Electoral Behavior*, eds. K. Arzheimer, J. Evans, and M. S. Lewis-Beck. SAGE, 584–605.
- Steinmetz, Katy. 2018. "The Real Fake News Crisis." *TIME Magazine* 192(7).
- Stelter, Jonathan. 2011. "Occupy Wall Street Puts the Coverage in the Spotlight - The New York Times." *The New York Times*.
<https://www.nytimes.com/2011/11/21/business/media/occupy-wall-street-puts-the-coverage-in-the-spotlight.html> (Accessed April 2nd, 2021).
- Storer v. Brown*. 1974. (415 U.S. 724).
- Strevens, Michael. 2003. "The Role of the Priority Rule in Science." *The Journal of Philosophy* 100(2): 55–79.
- Sunstein, Cass R. 2014. *On Rumors*. Princeton University Press.
- . 2017. *#Republic : Divided Democracy in the Age of Social Media*. Princeton: Princeton University Press.
- Sunstein, Cass R., and Adrian Vermeule. 2009. "Conspiracy Theories: Causes and Cures." *Journal of Political Philosophy* 17(2).
- "Syria Conflict: 'Chemical Attack' in Idlib Kills 58." 2017. *BBC News*.
<https://www.bbc.com/news/world-middle-east-39488539> (Accessed April 2nd, 2021).
- Taber, Charles S., and Milton Lodge. 2006. "Motivated Skepticism in the Evaluation of Political Beliefs." *American Journal of Political Science* 50(3): 755–69.
- Tankovska, H. 2021. "Social Media - Statistics & Facts." *Statista*.
<https://www.statista.com/topics/1164/social-networks/> (Accessed April 1st, 2021).
- Teovanović, Predrag, Goran Knežević, and Lazar Stankov. 2015. "Individual Differences in Cognitive Biases: Evidence against One-Factor Theory of Rationality." *Intelligence* 50: 75–86.
- The Vaccine War*. 2010. PBS. <https://www.pbs.org/wgbh/frontline/film/vaccines/>.

- “Thimerosal and Vaccines | Vaccine Safety | CDC.” 2020. *Centers for Disease Control and Prevention*. <https://www.cdc.gov/vaccinesafety/concerns/thimerosal/index.html> (Accessed April 2nd, 2021).
- Traub, James. 2016. “It’s Time for the Elites to Rise Up Against the Ignorant Masses.” *Foreign Policy*. <https://foreignpolicy.com/2016/06/28/its-time-for-the-elites-to-rise-up-against-ignorant-masses-trump-2016-brexit/> (Accessed April 2nd, 2021).
- Tucker, Heather. 2017. “AP: Larry Nassar Expected to Plead Guilty, Faces at Least 25 Years in Prison.” *USA Today*. <https://www.usatoday.com/story/sports/olympics/2017/11/21/larry-nassar-usa-gymnastics-expected-plead-guilty-sexual-assault/886906001/> (Accessed April 2nd, 2021).
- Turner, Stephen P. 2003. *Liberal Democracy 3.0 : Civil Society in an Age of Experts*. London: SAGE Publications.
- Tversky, A., and D. Kahneman. 1974. “Judgment under Uncertainty: Heuristics and Biases.” *Science* 185(4157): 1124–31.
- Ungar, Sheldon. 1992. “The Rise and (Relative) Decline of Global Warming as a Social Problem.” *The Sociological Quarterly* 33(4): 483–501.
- United States of America: WHO and UNICEF Estimates of Immunization Coverage: 2018 Revision*. 2018. World Health Organization.
- Van Schoelandt, Chad. 2015. “Justification, Coercion, and the Place of Public Reason.” *Philosophical Studies : An International Journal for Philosophy in the Analytic Tradition* 172(4): 1031–50.
- Van Swol, Lyn M., Deepak Malhotra, and Michael T. Braun. 2012. “Deception and Its Detection: Effects of Monetary Incentives and Personal Relationship History.(Author Abstract).” *Communication Research* 39(2): 217–38.
- Vanderslott, Samantha. 2019. “Exploring the Meaning of Pro-Vaccine Activism across Two Countries.” *Social Science & Medicine* 222: 59–66.
- Viehoff, Daniel. 2016. “Authority and Expertise.” *Journal of Political Philosophy* 24(4).
- Vosoughi S, Roy D, and Aral S. 2018. “The Spread of True and False News Online.” *Science* 359(6380): 1146–51.
- Waldron, Jeremy. 1987. “Theoretical Foundations of Liberalism.” *The Philosophical Quarterly* 37(147): 127–50.
- Wardle, Claire, and Hossein Derakhshan. 2017. *Information Disorder: Toward an Interdisciplinary Framework for Research and Policy Making*. Council of Europe.

- Warren, Mark (Mark E.). 1999. *Democracy and Trust*. Cambridge, UK ; Cambridge University Press.
- Waszak, Przemyslaw M., Wioleta Kasprzycka-Waszak, and Alicja Kubanek. 2018. "The Spread of Medical Fake News in Social Media - The Pilot Quantitative Study." *Health Policy and Technology* 7(2): 115–18.
- Weisberg M. and Muldoon R. 2009. "Epistemic Landscapes and the Division of Cognitive Labor." *Philosophy of Science* 76(2): 225–52.
- Wenar, Leif. 2010. "Poverty Is No Pond: Challenges for the Affluent." In *Giving Well: The Ethics of Philanthropy*, eds. T. Illingworth, T. Pogge, and L. Wenar. Oxford University Press, 104–32.
- "What Is the U.S. Opioid Epidemic?" 2017. *United States Department of Health and Human Services*. <https://www.hhs.gov/opioids/about-the-epidemic/index.html> (Accessed April 2nd, 2021).
- Whitsett, J. David. 2019. "Fake News According to Trump." Ohio State University.
- Williams, Casey. 2017. "Has Trump Stolen Philosophy's Critical Tools?" *The New York Times*.
- Wingfield, Nick, Mike Isaac, and Katie Benner. 2016. "Google and Facebook Take Aim at Fake News Sites." *The New York Times*. <https://www.nytimes.com/2016/11/15/technology/google-will-ban-websites-that-host-fake-news-from-using-its-ad-service.html> (Accessed April 2nd, 2021).
- Wodak, Ruth. 2019. "Entering the 'Post-Shame Era': The Rise of Illiberal Democracy, Populism and Neo-Authoritarianism in EUrope." *Global Discourse* 9(1): 195–213.
- "Worldnewsdailyreport.Com Traffic, Ranking & Marketing Analytics." 2021. *SimilarWeb*. <http://similarweb.com/website/worldnewsdailyreport.com/> (Accessed April 1st, 2021).
- Zagzebski, Linda Trinkaus 1946-. 2012. *Epistemic Authority : A Theory of Trust, Authority, and Autonomy in Belief*. Oxford: Oxford University Press.
- Zakaria, Fareed. 1997. "The Rise of Illiberal Democracy." *Foreign Affairs* 76(6): 22–43.
- Zaller, John, and Dennis Chiu. 1996. "Government's Little Helper: U.S. Press Coverage of Foreign Policy Crises, 1945–1991." *Political Communication* 13(4): 385–405.
- Zollman, Kevin J. S. 2010. "The Epistemic Benefit of Transient Diversity." *Erkenntnis* 72(1): 17–35.