

OPTIMAL FORAGING THEORY REVISITED

A Thesis

Presented in Partial Fulfillment of the Requirements for
the Degree Master of Science in the
Graduate School of The Ohio State University

By

Theodore P. Pavlic, B.S.

* * * * *

The Ohio State University

2007

Master's Examination Committee:

Kevin M. Passino, Adviser

Yuan F. Zheng

Thomas A. Waite

Approved by

Adviser

Electrical & Computer
Engineering Graduate
Program

© Copyright by
Theodore P. Pavlic
2007

ABSTRACT

Optimal foraging theory explains adaptation via natural selection through quantitative models. Behaviors that are most likely to be favored by natural selection can be predicted by maximizing functions representing Darwinian fitness. Optimization has natural applications in engineering, and so this approach can also be used to design behaviors of engineered agents. In this thesis, we generalize ideas from optimal foraging theory to allow for its easy application to engineering design. By extending standard models and suggesting new value functions of interest, we enhance the analytical efficacy of optimal foraging theory and suggest possible optimality reasons for previously unexplained behaviors observed in nature. Finally, we develop a procedure for maximizing a class of optimization functions relevant to our general model. As designing strategies to maximize returns in a stochastic environment is effectively an optimal portfolio problem, our methods are influenced by results from modern and post-modern portfolio theory. We suggest that optimal foraging theory could benefit by injecting updated concepts from these economic areas.

This is dedicated to my brother Kenny, whose bright disposition in dark times is not only illuminating but warming. I could not be more proud to be a part of a family that could produce someone like him.

ACKNOWLEDGMENTS

First, I give thanks to my parents, Paul and Eileen, who have always been supportive and understanding, even when research has reduced the frequency of contact with them. Any success that I have today could not have been possible without them.

I am also thankful for my girlfriend Jessie, who not only has tolerated long work nights but has also managed to prevent me from starvation. I value her support and encouragement. My hope to maximize the time I spend with her has been a strong impetus to proceed quickly in my research.

My adviser, Professor Kevin M. Passino, deserves thanks not only for his wisdom and guidance but also for his unending patience with me. Through him, I have not only learned engineering, but I have become a better writer and overall thinker. He has strengthened my understanding of how to research effectively and continues to serve as an important role-model for me.

Any accurate understanding that I have of behavioral ecology is entirely due to Professor Thomas A. Waite. The tangible and intangible benefits of collaboration with him are too numerous to list. I am an interloper in his field, and he has not only tolerated my intrusion but has welcomed me and provided me with instructions on how I might proceed deeper into new spaces. Exposure to him and his colleagues has left me in awe of the ecological adventures that are common in his field.

I also owe thanks to Professor Jerry F. Downhower for teaching me about natural selection. His teachings have attenuated my ignorance and improved my ability to understand not only the language of biology but also the world around me. I am sure that all of his students must feel the same way.

Real analysis and the study of stochastic processes have been the two most useful tools that I use regularly my graduate work. The sophistication I have gained in the former is entirely due to Professor Neil Falkner, whose attention to detail is admirable. My understanding of the latter comes from Professor Randolph L. Moses, who is one of the best teachers that I know. I am grateful for the time that both have volunteered to answering questions of mine. The expertise of Professor Jose B. Cruz Jr. has also been invaluable to me. His teachings about optimization have influenced much of the content of this thesis.

I must thank Professor James N. Bodurtha Jr. for introducing me to the field of finance and being patient with my elementary questions. I wish I could have spent more time investigating the fascinating problems and results studied in this area. The little that I have found with his guidance has been extremely useful and suggests to me many future research directions.

While I have never met Matt Ridley, long ago his popularizations of human behavior and genetics are what encouraged me to learn about biology. Therefore, I owe him thanks. Without his fascinating works, this thesis would most likely be far more conventional.

Finally, I thank Professor Yuan F. Zheng for agreeing to take the time to be a member of my thesis committee. As he is an expert in biological applications of robotics, I am eager to hear his responses to my work.

VITA

February 28, 1981	Born - Columbus, OH, USA
June 2004	B.S., Elec. & Comp. Engineering
2004–present	Dean’s Distinguished Univ. Fellow, The Ohio State University
2006–2007	NSF GK-12 Fellow, The Ohio State University
2002, 2003	Analog Design Intern, National Instruments, Austin, Texas
2001	Core Systems Developer, IBM Storage, RTP, North Carolina

PUBLICATIONS

Research Publications

R. J. Freuler, M. J. Hoffmann, T. P. Pavlic, J. M. Beams, J. P. Radigan, P. K. Dutta, J. T. Demel, and E. D. Justen. Experiences with a comprehensive freshman hands-on course – designing, building, and testing small autonomous robots. In *Proceedings of the 2003 American Society for Engineering Education Annual Conference & Exposition*, 2003.

T. P. Pavlic and K. M. Passino. Foraging theory for mobile agent speed choice. *Engineering Applications of Artificial Intelligence*. Submitted.

FIELDS OF STUDY

Major Field: Electrical & Computer Engineering

TABLE OF CONTENTS

	Page
Abstract	iii
Dedication	iv
Acknowledgments	v
Vita	vii
List of Tables	xi
List of Figures	xii
Chapters:	
1. Introduction	1
2. Model of a Solitary Agent	3
2.1 The Generalized Solitary Agent Model	6
2.1.1 Model Assumptions	6
2.1.2 Task-Type Parameters	8
2.1.3 Actual Processing Gains, Costs, and Times	12
2.1.4 Important Technical Notes	12
2.2 Classical OFT Analysis: Encounter-Based Approach	13
2.2.1 Processes Generated from Merged Encounters	13
2.2.2 Markov Renewal Process	16
2.2.3 Markov Renewal-Reward Processes	18
2.2.4 Reward Process Statistics	18
2.3 Finite Lifetime Analysis: Processing-Based Approach	22
2.3.1 Poisson Encounters of Processed Tasks of One Type	22

2.3.2	Process-Only Markov Renewal Process	24
2.4	Relationship Between Analysis Approaches	29
2.5	Weaknesses of the Model	30
3.	Statistical Optimization Objectives for Solitary Behavior	33
3.1	Objective Function Structure	34
3.1.1	Statistics of Interest	34
3.1.2	Optimization Constraints	36
3.1.3	Impact of Function Choice on Optimal Behaviors	38
3.2	Classical OFT Approach to Optimization	39
3.2.1	Maximization of Long-Term Rate of Net Gain	39
3.2.2	Minimization of Net Gain Shortfall	44
3.2.3	Criticisms of the OFT Approach	48
3.3	Generalized Optimization of Solitary Agent Behavior	49
3.3.1	Finite Task Processing	50
3.3.2	Tradeoffs as Ratios	51
3.3.3	Generalized Pareto Tradeoffs	63
3.3.4	Constraints	66
3.4	Future Directions Inspired by PMPT	70
3.4.1	Lower Partial Moments	70
3.4.2	Stochastic Dominance	72
4.	Finite-Lifetime Optimization Results	74
4.1	Optimization of a Rational Objective Function	74
4.1.1	The Generalized Problem	75
4.1.2	The Optimization Procedure	76
4.1.3	Solutions to Special Cases	84
4.2	Optimization of Specific Objective Functions	90
4.2.1	Maximization of Rate of Excess Net Point Gain	90
4.2.2	Maximization of Discounted Net Gain	91
4.2.3	Maximization of Rate of Excess Efficiency	92
5.	Conclusion	93
5.1	Contributions to Engineering	93
5.2	Contributions to Biology	94
5.3	Future Directions	96
5.4	The Value of Collaboration	97

Appendices:

A. Limits of Markov Renewal Processes	98
Bibliography	100
List of Acronyms	106
List of Terms	107
List of Symbols	109
Index	114
People	121

LIST OF TABLES

Table	Page
3.1 Common Statistics for Solitary Optimization	35

LIST OF FIGURES

Figure	Page
2.1 Classical OFT Markov Renewal Process	16
2.2 Process-Only Markov Renewal Process	25
3.1 Visualization of Classical OFT Rate Maximization	43
3.2 Visualization of Classical OFT Risk-Sensitive Solutions	48
3.3 Visualization of Rate Maximization	53
3.4 Visualization of Efficiency Maximization	56
3.5 Visualization of Reward-to-Variability Maximization	60
3.6 Visualization of Reward-to-Variance Maximization	62

CHAPTER 1

INTRODUCTION

Following the example of Andrews et al. [1], Andrews et al. [2], Pavlic and Passino [46], and Quijano et al. [50], we synthesize ideas from Stephens and Krebs [60] to apply optimal foraging theory (OFT) to engineering applications. In particular, we expand the solitary agent framework from classical OFT so that it applies to more general cases. This framework describes a solitary agent (e.g., an autonomous vehicle) that faces tasks to process at random. On encounters with a task, the designed agent behavior specifies whether or not the agent should process the task and for how long processing should continue. This is inherently an optimal portfolio [36] problem as it involves allocating resources (e.g., time and cost of processing) in a way that optimizes some aspect of random future returns (e.g., value of tasks relative to fuel cost). Therefore, we then derive optimization results in this framework using methods borrowed from optimal portfolio theory. We hope that these extensions of OFT will be useful in the design of high-level control of autonomous agents and will also provide new insights in biological applications.

In [Chapter 2](#), we use insights from behavioral ecology to develop a general stochastic model of a solitary agent with statistics that may be used in analyzing or designing optimal behavior. In particular, we generalize the stochastic model used by classical

OFT and propose a new analysis approach. The statistics used in classical OFT are conditioned on the number of tasks *encountered* regardless of whether or not those tasks are processed. In our approach, we focus on statistics conditioned on the number of tasks *processed*. Not only does this have greater applicability to engineering, but it provides a new method for finite-lifetime analysis.

In [Chapter 3](#), we study various ways that statistics of our generalized agent may be combined for multiobjective optimization. We first describe the approaches used in classical OFT. By generalizing these classical objectives, we suggest new explanations for peculiar foraging behaviors observed in nature. We then propose new optimization objectives for use in engineering; however, we discuss how these objectives may also be applicable in behavioral ecology. Finally, we discuss how existing work in classical OFT may be duplicating existing work in economics. We suggest that a study of the most recent optimal portfolio theory literature may provide valuable insights to both behavioral analysis and design.

In [Chapter 4](#), we analyze a class of optimization functions that share a particular structure. Many of the functions we introduce in [Chapter 3](#) for multiobjective optimization have this structure, and so this analysis leads to optimal solutions for them. We present some of those solutions at the end of the chapter.

Concluding remarks are given in [Chapter 5](#). [Appendix A](#) provides some results from renewal theory that are used in [Chapter 2](#). Lists of acronyms, model terms, and mathematical symbols that we use are given at the end of this document. Topic and people indices follow the bibliography.

CHAPTER 2

MODEL OF A SOLITARY AGENT

In this chapter, we present a stochastic model of a typical solitary agent (i.e., neither competition nor cooperation is modeled) as a generalization of the one described by Charnov and Orians [16]. This model is similar to numerous deterministic and stochastic foraging models in the ecology literature [e.g., 14, 15, 25, 47, 48, 55, 67]; we focus on the model of Charnov and Orians because its high level of mathematical rigor lets it encompass many features of most other models in a theoretically convincing way. Introducing additional generality to this model allows it to be used in a wider range of applications that have different optimization criteria than classical OFT. We also suggest a new way of deriving statistics for this model based on a fixed number of tasks *processed*. This differs from the conventional statistical approach in OFT which focusses on statistics based on a fixed number of tasks *encountered* regardless of processing. Our approach has wider application to engineering and provides a new way of handling analysis of finite-lifetime behavior.

Below, we introduce terminology that will be used throughout this document and give the motivations for our approach. The model is presented in Section 2.1. In Section 2.2, we describe the analytical approach used in classical OFT. We present our approach as a modification to the classical OFT method in Section 2.3. Interesting

relationships between the two methods are given in [Section 2.4](#). Finally, weaknesses of this model (and thus also of both approaches) are given in [Section 2.5](#). A list of some frequently used terms in this model and the two approaches is given [at the end of this document](#).

Terminology: Agents, Tasks, and Currency

The model we use describes a generic *agent* that *searches* at some constant rate for *tasks* to *process* in an effort to acquire *point gain*. The agent is assumed to be able to detect all potential tasks perfectly. During both searching and processing, the agent may have to pay *costs*; however, the agent will pay no cost to detect the tasks. The point gain and costs will be given in the same *currency*, and so *net point gain* will be the difference between point gain and costs. For example, this model could describe an animal foraging for energetic gain at some energetic cost, or it could describe an autonomous military vehicle searching for targets at the expense of fuel.

Behavioral Optimization: Making the Best Choices

When an agent encounters a task, we refer to making a *choice* among different behavioral options within the model for processing that task. Despite this naming convention, we do not imply that the agent needs to have the cognitive ability to make choices; the agent only needs to behave in some consistent manner. We then can build performance measures over the space of these behaviors. In a biological context, these performance measures may model reproductive success. In an engineering context, these performance measures may, for example, measure the relative importance of various tasks with respect to the fuel cost required to complete them.

Whether through natural selection or engineering design, behaviors that optimize these performance measures should be favored.

Approach Motivation: Finite Lifetime Analysis and Design

Our model is more than just semantically different than the classical OFT model originally introduced by Charnov and Orians [16] and popularized by Stephens and Krebs [60]. For one, it takes parameters from a wider range of values and replaces deterministic aspects of the OFT model with first-order statistics of random variables. More importantly, our new approach to analysis provides a convenient method for analyzing behavior over a *finite* lifetime (or *runtime* in an engineering context). Classical OFT does not attempt to analyze finite lifetimes. Instead, limiting statistics on a space of never-ending behaviors are used. It is natural to define a finite lifetime as a finite number of tasks processed. However, classical OFT focusses its analysis on cycles that start and end on task encounters regardless of whether those encounters lead to processing. In our approach, we recognize that because the agent does not pay a recognition cost on each encounter, all encounters that do not result in processing may be discarded. Because we consider only the encounters that result in processing, a finite lifetime can be defined as a finite number of these encounters. This can be useful, for example, if processing a task involves depositing one of a limited number of objects.

2.1 The Generalized Solitary Agent Model

An agent's lifetime is a random experiment modeled by the probability space¹ $(\mathcal{U}, \mathcal{P}(\mathcal{U}), \text{Pr})$. That is, each outcome $\zeta \in \mathcal{U}$ represents one possible lifetime for the agent, and so we will often substitute the term *lifetime* for the term *outcome*. Thus, statistics on random variables² in this probability space will include parameters that fully specify the environment and the agent's behavior. For example, if the agent acquires gain over its lifetime, the expected³ gain represents the probabilistic average of all possible gains given the agent's behavior and the randomness in the environment. The optimization goal will be to choose behavioral parameters that yield the optimum statistics in the given environment.

2.1.1 Model Assumptions

An agent's lifetime (i.e., each random outcome in the model) consists of searching for tasks, choosing whether to process those tasks, processing those tasks, receiving gains for processing those tasks, and paying costs for searching and processing. The following are general assumptions about these aspects of the agent's interaction with its environment.

Independent Processing Cost Rates: Processing costs are linear in processing time, and so they are completely specified by *processing cost rates*. We assume these

¹A *probability space* is a set of outcomes, a set of events that each are a set of outcomes, and a measure mapping those events to their probability.

²A *random variable* X is a measurable function mapping events into Borel sets of real numbers.

³The *expectation* $E(X)$ is $\int_{-\infty}^{\infty} x f_X(x) dx$ where f_X is the (Lebesgue) probability density of events under X . The expectation is often called the *mean* or the *(first) moment (about the origin)*. It represents the *center* of mass of the distribution.

cost rates are uncorrelated⁴ with any length of (processing) time, and that the processing cost of any particular task is independent⁵ of the processing cost of any other task.

Independent Processing Gains: The processing gain for any particular task is independent of the processing gain of any other task.

Independent Processing Decisions: An agent's decision to process any particular task is independent of its decision to process any other task.

Pseudo-Deterministic Search Cost Rate: The search cost for finding any particular task is assumed to be independent of the type of that task and independent of the search cost of finding any other task. Additionally, search costs are assumed to be linear in search time, and so they are completely specified by *search cost rates*. We make several assumptions about these rates.

- Search cost rates are uncorrelated with any length of time.
- For any lifetime $\zeta \in \mathcal{U}$, the search cost rate is a single random variable rather than some kind of random process. In other words, we assume the search cost rate is constant over the entire lifetime of an agent. Thus, we consider the search cost rate to be the random variable $C^s : \mathcal{U} \mapsto \overline{\mathbb{R}}$.
- We define $c^s \in \mathbb{R}$ as the expectation of random variable C^s (i.e., $c^s = E(C^s)$), so c^s is finite.

⁴To say random variables X and Y are uncorrelated means $E(XY) = E(X)E(Y)$.

⁵To say random variables X , Y , and Z are (mutually) independent means that $f_{XYZ}(x, y, z) = f_X(x)f_Y(y)f_Z(z)$. This implies that they are uncorrelated and that $E(X|Y) = E(X)$.

- We assume $\Pr(C^s = c^s) = 1$. This is roughly equivalent to assuming that C^s is deterministic. This assumption is critical for the analyses of variance and stochastic limits in the model; if neither of these is of interest, then this assumption can be relaxed entirely.

Thus, in many cases, the parameter c^s will be an acceptable surrogate for the phrase *search cost rate* or even *search cost* as long as it is understood to be a rate.

2.1.2 Task-Type Parameters

Tasks encountered by an agent during its lifetime are grouped into types that share certain characteristics. In particular, there are $n \in \mathbb{N}$ distinct task types. Take $i \in \{1, 2, \dots, n\}$.

Task-Type Processes: For task type i , encounters are driven by a Poisson process $(M_i(t_s) : t_s \in \mathbb{R}_{\geq 0})$. That is, for each lifetime $\zeta \in \mathcal{U}$, $M_i(t_s)$ is the number of encounters with tasks of type i after $t_s \in \mathbb{R}_{\geq 0}$ units of *search time*. We associate the following sequences of (mutually) independent and identically distributed (i.i.d.) random variables with finite expectation⁶ with this Poisson process.

- (I_M^i) : Random process representing the type of the task. That is, $I_M^i = i$ for all $N \in \mathbb{N}$ and all $\zeta \in \mathcal{U}$.
- (g_M^i) : Random process representing potential gross processing gains (i.e., the gross gain rewarded if the task is chosen for processing) for encounters with tasks of type i .

⁶To say random variable X has finite expectation means that $E(|X|) < \infty$.

- (τ_M^i) : Random process representing potential processing times (i.e., the processing time if the task is chosen for processing) for encounters with tasks of type i .
- (c_M^i) : Random process representing potential cost rates (i.e., the cost rate for processing time if the task is chosen for processing) for encounters with tasks of type i . Thus, $(c_M^i T_M^i)$ is a random process of potential costs (i.e., the processing cost if the task is chosen for processing) for encounters with tasks of type i .
- (X_M^i) : Random process representing the agent's choice to process a task of type i immediately after encountering it. That is, for encounter $N \in \mathbb{N}$ of lifetime $\zeta \in \mathcal{U}$,

$$X_M^i = \begin{cases} 0 & \text{if the agent chooses not to process the task} \\ 1 & \text{if the agent chooses to process the task} \end{cases}$$

We make several assumptions about this process.

- (i) For all $N \in \mathbb{N}$, $\mathbb{E}(X_M^i) = 1$ if and only if $X_M^i(\zeta) = 1$ for all $\zeta \in \mathcal{U}$.
- (ii) For all $N \in \mathbb{N}$, $\mathbb{E}(X_M^i) = 0$ if and only if $X_M^i(\zeta) = 0$ for all $\zeta \in \mathcal{U}$.
- (iii) Each processing choice is independent of all other processing choices.
- (iv) For $M \in \mathbb{N}$, X_M is uncorrelated with $(g_M^i - c_M^i \tau_M^i)$, $c_M^i \tau_M^i$, and τ_M^i .

It is clear that (X_M^i) is a sequence of Bernoulli trials.

Parameters of Task Types: The above random processes are characterized by the parameters below. Tasks within a particular type all share these parameters; that is, these parameters also characterize each task type.

- $\lambda_i \in \mathbb{R}_{>0}$: The Poisson rate for process $(M_i(t_s) : t_s \in \mathbb{R}_{\geq 0})$ (i.e., $\lambda_i = 1/\mathbb{E}(T_1^i)$). An expanded version of this model might introduce detection errors by modulating this parameter, which might also be made to depend on search speed. Pavlic and Passino [46] incorporate both of these aspects with the analogous parameter of a similar agent model.
- $\tau_i \in \mathbb{R}$: The average processing time, given in seconds, for processing a task of type i (i.e., $g_i = \mathbb{E}(\tau_1^i)$).
- $c_i \in \mathbb{R}$: The average fuel cost rate, given in points per second, for processing a task of type i (i.e., $c_i = \mathbb{E}(c_1^i)$).
- $g_i \in \mathbb{R}$: The average gross gain, given in points, for processing a task of type i (i.e., $g_i = \mathbb{E}(g_1^i)$).
- $p_i \in [0, 1]$: An agent's preference for processing a task of type i .
 - If $p_i = 0$, then no tasks of type i are processed.
 - If $p_i \in (0, 1)$, then tasks of type i are processed according to successes of a Bernoulli trial with parameter p_i .
 - If $p_i = 1$, then all tasks of type i are processed.

That is, p_i can be called the probability that the agent will process a task of type i (i.e., $\mathbb{E}(X_1^i) = p_i$). Detection errors could be introduced via this parameter as well.

Of course, it is trivial that $\mathbb{E}(I_1^i) = i$.

Average Gain as Function of Average Time: Unlike with processing costs, the relationship between processing time and processing gain has not been made

explicit. In general, the model of the system will require g_i to change whenever τ_i changes. That is, it makes sense that a longer average processing time would alter the average gain. Therefore, we introduce the function $g_i : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}$ so that $g_i(\tau_i)$ represents the average gain returned from tasks of type i given an average processing length of $\tau_i \in \mathbb{R}_{\geq 0}$. This function is used when predicting the optimal processing time in a given environment. We usually assume g_i is continuously differentiable.

Optimization Variables and Prey and Patch Models: The behavior of an agent is completely specified by the preference probabilities (i.e., p_i for all $i \in \{1, 2, \dots, n\}$) and the processing times (i.e., τ_i for all $i \in \{1, 2, \dots, n\}$). All other parameters are fixed with the agent's environment. The *task processing-length choice problem* refers to the case when the preference probabilities are also fixed with the environment (i.e., absorbed into the task type encounter rates) so that the agent is free to choose processing times only; this is called a *patch model* by biologists [60]. The *task-type choice problem* refers to the case when the processing times are fixed with the environment so that the agent is free to choose preference probabilities only; this is called a *prey model* by biologists [60]. The most general case, when the agent is free to choose both, is called the *combined task-type and processing-length choice problem*; biologists refer to this case as the *combined prey and patch model* [60].

These processes and parameters will be used throughout this document.

2.1.3 Actual Processing Gains, Costs, and Times

Take $i \in \{1, 2, \dots, n\}$. For the rest of this chapter, we will also use the processes (G_M^i) , (C_M^i) , and (T_M^i) , which are defined with $G_M^i \triangleq X_M^i g_M^i$ and $C_M^i \triangleq X_M^i c_M^i \tau_M^i$ and $T_M^i \triangleq X_M^i \tau_M^i$ for all $\zeta \in \mathcal{U}$ and $N \in \mathbb{N}$. These represent the actual processing gain, processing cost, and processing time for each task encounter. Clearly, the gain (processing time) of any task is independent of the gain (processing time) of any other task; additionally, (G_M^i) , (C_M^i) , and (T_M^i) are sequences of i.i.d. random variables with finite expectation. It is necessary for $\Pr(C^s = c^s) = 1$ for the random variables of (G_M^i) and (C_M^i) to be i.i.d.. If this is not the case, then the random variables of (G_M^i) and (C_M^i) will be identically but not independently distributed.

2.1.4 Important Technical Notes

This model has more flexibility than the classical OFT models described by Stephens and Krebs [60]. It also shares one aspect of classical OFT foraging models that is often taken for granted.

Enhanced Gain and Cost Structure: We augment the conventional classical OFT foraging model with time-dependent costs, while not restricting the signs of our cost and gains. That is, we allow costs and gains to be positive, zero, or negative. In other words, negative costs may be viewed as time-dependent gains just as negative gains may be viewed as time-constant costs. For example, a negative search cost may be viewed as modeling the value of some other useful activity that can only be done during searching. Some impacts of this generalization of the gain and cost structure are discussed in [Chapter 3](#).

Poisson Processes and Simultaneous Encounters: All of the assumptions listed in [Sections 2.1.1](#) and [2.1.2](#) are important, but one particular assumption (that is also found in the classical solitary foraging model) deserves special attention, namely that model encounters occur according to a Poisson process. A consequence of this assumption is that interarrival times have a particular continuous distribution. Additionally, this assumption implies that simultaneous encounters occur with probability zero; therefore, behavioral statistics are not affected by the choices made by the agent on a simultaneous encounter.

2.2 Classical OFT Analysis: Encounter-Based Approach

Here, we introduce an approach to analysis of agent behavior based on classical OFT [e.g., [16](#), [60](#)]. We call this a *merge before split* approach. In this approach, the encounter rates of each type are independent of the preference probabilities. That is, the agent is considered to encounter each task and then choose whether to process the task. Because encounters are generated by Poisson processes, an alternative approach would be to make the preference probabilities a modifier of the encounter rates rather than some aspect of the agent's choice; this alternative is described in [Section 2.3](#). The merged processes generated by encounters with all tasks are described in [Section 2.2.1](#). [Sections 2.2.2](#), [2.2.3](#), and [2.2.4](#) use renewal theory based on these merged processes to develop statistics that can be used as optimization criteria for agent behavior.

2.2.1 Processes Generated from Merged Encounters

Above, we defined n Poisson processes corresponding to the n task types. However, as an agent searches, it encounters tasks from n processes at once. That is, the agent

faces the merged Poisson process $(M(t_s) : t_s \in \mathbb{R}_{\geq 0})$ defined for all $\zeta \in \mathcal{U}$ and all $t_s \in \mathbb{R}_{\geq 0}$ by

$$M(t_s) \triangleq \sum_{i=1}^n M_i(t_s)$$

which carries with it the interevent time process (Υ_M) . In other words, for any lifetime $\zeta \in \mathcal{U}$, $M^p(t_s)$ represents the number of tasks *encountered* after *searching* for t_s time. We call the encounter rate for this process λ , where $\lambda = \sum_{i=1}^n \lambda_i$ by the theory of merged Poisson processes [64]. Therefore, $E(\Upsilon_1) = 1/\lambda$. Because this process is also a Markov renewal process, $\text{aslim}_{t_s \rightarrow \infty} M(t_s) = \infty$; however, because this is a Poisson counting process, $E(M(t_s)) = \lambda t_s$ for all $t_s \in \mathbb{R}_{\geq 0}$.

Merged Task-Type Processes

Define the random processes (\mathcal{D}_M) , (\mathcal{U}_M) , (\mathcal{T}_M) , and (I_M) as merged versions of the families $((G_M^i)_{i=1}^n)$, $((C_M^i)_{i=1}^n)$, $((T_M^i)_{i=1}^n)$, and $((I_M^i)_{i=1}^n)$ respectively. Each of these processes is an i.i.d. sequence of random variables. The random variables I_1 and Υ_1 are assumed to be independent. For any lifetime $\zeta \in \mathcal{U}$, $I_1 = i$ would indicate that the first encounter was generated by process $(M_i(t_s) : t_s \in \mathbb{R}_{\geq 0})$. It will be convenient for us to introduce the symbols g , c , and τ defined by

$$g \triangleq E(\mathcal{D}_1) \quad \text{and} \quad c \triangleq E(\mathcal{U}_1) \quad \text{and} \quad \tau \triangleq E(\mathcal{T}_1)$$

These random variables respectively represent the net gain, cost, and time for processing a task during a single arbitrary OFT renewal cycle. We also use the notation \bar{g} , \bar{c} , and $\bar{\tau}$ defined by

$$\bar{g} \triangleq E(g) = E(\mathcal{D}_1) \quad \text{and} \quad \bar{c} \triangleq E(c) = E(\mathcal{U}_1) \quad \text{and} \quad \bar{\tau} \triangleq E(\tau) = E(\mathcal{T}_1)$$

From the theory of merged Poisson processes, $\Pr(I_1 = i) = \lambda_i/\lambda$ for all $i \in \{1, 2, \dots, n\}$.

Combining this with the fact that $\lambda = \sum_{i=1}^n \lambda_i$ and a property⁷ of expectation yields

$$\bar{g} = \sum_{j=1}^n \frac{\lambda_j}{\lambda} p_j g_j \quad \text{and} \quad \bar{c} = \sum_{j=1}^n \frac{\lambda_j}{\lambda} p_j c_j \tau_j \quad \text{and} \quad \bar{\tau} = \sum_{j=1}^n \frac{\lambda_j}{\lambda} p_j \tau_j$$

So, these expectations are weighted sums of parameters. In particular, if $n = 1$,

$$\bar{g} = p_1 g_1(\tau_1) \quad \text{and} \quad \bar{c} = p_1 c_1 \tau_1 \quad \text{and} \quad \bar{\tau} = p_1 \tau_1$$

This result is useful when visualizing optimization results. Additionally,

$$\mathbb{E}(C^s \Upsilon_1 | I_1 = i) = \mathbb{E}(C^s \Upsilon_1) = \frac{c^s}{\lambda}$$

Below, we use these results frequently in expressions of statistics.

Net Gain, Cost, and Time Processes

Now, we define random processes (\tilde{G}_N) , (\tilde{C}_N) , and (\tilde{T}_N) with

$$\tilde{G}_N \triangleq \mathfrak{D}_N - \mathfrak{U}_N - C^s \Upsilon_N \quad \text{and} \quad \tilde{C}_N \triangleq \mathfrak{U}_N + C^s \Upsilon_N \quad \text{and} \quad \tilde{T}_N \triangleq \mathfrak{T}_N + \Upsilon_N$$

for all $N \in \mathbb{N}$ and $\zeta \in \mathcal{U}$. It is clear that (\tilde{G}_N) , (\tilde{C}_N) , and (\tilde{T}_N) are i.i.d. sequences of random variables with finite expectation. In some cases, it will be interesting to look at the *gross gain* returned to an agent. Thus, we define the process $(\tilde{G}_N + \tilde{C}_N)$ as well⁸. By the above definitions, $\tilde{G}_1 + \tilde{C}_1 = g$ and $\tilde{G}_N + \tilde{C}_N = \mathfrak{D}_N$ for all $N \in \mathbb{N}$ and $\zeta \in \mathcal{U}$. The statistics of these random variables are of interest to us. In particular,

⁷For random variables X and Y , $\mathbb{E}(X) = \mathbb{E}(\mathbb{E}(X|Y))$.

⁸Recall that the all cost rates may be negative in this model. While these costs would be interpreted as gains in this case, they are not included in this definition of gross gain. Gross gain is all gains before the impact of costs, positive or negative.




Figure 2.1: The classical OFT Markov renewal process, where the solid dot is the renewal point that starts each cycle.

$$\mathbb{E}(\tilde{G}_1) = \bar{g} - \bar{c} - \frac{c^s}{\lambda} \quad (2.1)$$

$$\mathbb{E}(\tilde{C}_1) = \bar{c} + \frac{c^s}{\lambda} \quad (2.2)$$

$$\mathbb{E}(\tilde{T}_1) = \bar{\tau} + \frac{1}{\lambda} \quad (2.3)$$

$$\mathbb{E}(\tilde{G}_1 + \tilde{C}_1) = \bar{g} \quad (2.4)$$

Also, $\Pr(\tilde{T}_1 = 0) = 0$ because $\mathbb{E}(\tilde{T}_1) > 0$ and $\Pr(\Upsilon_1 = 0) = 0$

2.2.2 Markov Renewal Process

Because (\tilde{T}_N) is an i.i.d. sequence of random variables with $0 < \mathbb{E}(\tilde{T}_1) < \infty$ and $\Pr(\tilde{T}_1) = 0$, the process $(N(t) : t \in \mathbb{R}_{\geq 0})$ defined by

$$N(t) \triangleq \sup \left\{ N \in \mathbb{N} : \sum_{i=1}^N \tilde{T}_i \leq t \right\} = \sup \left\{ N \in \mathbb{N} : \sum_{i=1}^N (\Upsilon_i + \Upsilon_i) \leq t \right\}$$

for all $t \in \mathbb{R}_{\geq 0}$ and all $\zeta \in \mathcal{U}$ is a Markov renewal process with interarrival process (\tilde{T}_N) . This process represents the number of tasks *encountered* from time 0 to time t (i.e., t is a measure of the agent's lifetime, not how long the agent has searched).

This Markov renewal process is depicted in [Figure 2.1](#), and one iteration around this process will be known as an *OFT cycle*. That is, because the agent can choose to process or ignore a task, the holding time for the renewal process always includes some

search time and may include processing time if an encounter is followed by a decision to process the task. By definition of this process, simultaneous encounters occur with probability zero. As with any Markov renewal process, $\text{aslim}_{t \rightarrow \infty} N(t) = \infty$; however, while $E(M(t_s))$ is known for all $t_s \in \mathbb{R}_{\geq 0}$, a derivation of $E(N(t))$ for all $t \in \mathbb{R}_{\geq 0}$ is outside the scope of this work. Fortunately, applications rarely require the precise form of this expectation. Additionally, it is known that for all $\zeta \in \mathcal{U}$ and all $t \in \mathbb{R}_{\geq 0}$, $N(t) \leq M(t)$; therefore, $0 \leq E(N(t)) \leq \lambda t$ for all $t \in \mathbb{R}_{\geq 0}$.

Encounter Times: Statistics and Stochastic Limits

The process (\tilde{T}^N) defined with $\tilde{T}^N \triangleq \sum_{i=1}^N \tilde{T}_i$ for all $N \in \mathbb{N}$ and all $\zeta \in \mathcal{U}$ is the sequence of encounter times for $(N(t) : t \in \mathbb{R}_{\geq 0})$. Because (\tilde{T}_N) is an i.i.d. sequence of random variables with finite expectation,

$$E(\tilde{T}^N) = N E(\tilde{T}_1) = \frac{N}{\lambda} + N\bar{\tau}$$

for all $N \in \mathbb{N}$. It can be shown⁹ that

$$\text{aslim}_{t \rightarrow \infty} \frac{N(t)}{t} = \lim_{t \rightarrow \infty} \frac{E(N(t))}{t} = \text{aslim}_{N \rightarrow \infty} \frac{N}{\tilde{T}^N} = \lim_{N \rightarrow \infty} E\left(\frac{N}{\tilde{T}^N}\right) = \frac{1}{E(\tilde{T}_1)} \quad (2.5)$$

Therefore, the ratio $1/E(\tilde{T}_1)$ may be called the *long-term encounter rate* of $(N(t) : t \in \mathbb{R}_{\geq 0})$. Similarly, it is also the case that

$$\text{aslim}_{t \rightarrow \infty} \frac{\tilde{T}(t)}{t} = \lim_{t \rightarrow \infty} \frac{E(\tilde{T}(t))}{t} = 1$$

which is not surprising; that is, as the agent's lifetime increases, the time spent waiting for the very next task encounter becomes negligible.

⁹See Appendix A.

2.2.3 Markov Renewal-Reward Processes

The processes (\tilde{G}_N) and (\tilde{C}_N) can be viewed as sequences of gains and losses, respectively, corresponding to each $(N(t) : t \in \mathbb{R}_{\geq 0})$ encounter. Define the corresponding cumulative processes¹⁰ (\tilde{G}^N) , (\tilde{C}^N) , and $(\tilde{G}^N + \tilde{C}^N)$ with

$$\tilde{G}^N \triangleq \sum_{i=1}^N \tilde{G}_i \quad \text{and} \quad \tilde{C}^N \triangleq \sum_{i=1}^N \tilde{C}_i \quad \text{and} \quad \tilde{G}^N + \tilde{C}^N = \sum_{i=1}^N (\tilde{G}_i + \tilde{C}_i)$$

for all $N \in \mathbb{N}$ and all $\zeta \in \mathcal{U}$. Also define the Markov renewal-reward processes¹¹

$(\tilde{G}(t) : t \in \mathbb{R}_{\geq 0})$, $(\tilde{C}(t) : t \in \mathbb{R}_{\geq 0})$, and $(\tilde{T}(t) : t \in \mathbb{R}_{\geq 0})$ with

$$\tilde{G}(t) \triangleq \tilde{G}^{N(t)} = \sum_{i=1}^{N(t)} \tilde{G}_i \quad \text{and} \quad \tilde{C}(t) \triangleq \tilde{C}^{N(t)} = \sum_{i=1}^{N(t)} \tilde{C}_i \quad \text{and} \quad \tilde{T}(t) \triangleq \tilde{T}^{N(t)} = \sum_{i=1}^{N(t)} \tilde{T}_i$$

and the process $(\tilde{G}(t) + \tilde{C}(t) : t \in \mathbb{R}_{\geq 0})$ accordingly with

$$\tilde{G}(t) + \tilde{C}(t) = \tilde{G}^{N(t)} + \tilde{C}^{N(t)} = \sum_{i=1}^{N(t)} (\tilde{G}_i + \tilde{C}_i)$$

for all $t \in \mathbb{R}_{\geq 0}$ and $\zeta \in \mathcal{U}$.

2.2.4 Reward Process Statistics

Because (\tilde{G}_N) and (\tilde{C}_N) are i.i.d. sequences of random variables with finite expectation, for all $N \in \mathbb{N}$,

$$\mathbb{E}(\tilde{G}^N) = N \mathbb{E}(\tilde{G}_1) = N \left(\bar{g} - \bar{c} - \frac{c^s}{\lambda} \right) \quad (2.6)$$

$$\mathbb{E}(\tilde{C}^N) = N \mathbb{E}(\tilde{C}_1) = N \left(\bar{c} + \frac{c^s}{\lambda} \right) \quad (2.7)$$

and, as we showed above,

$$\mathbb{E}(\tilde{T}^N) = N \mathbb{E}(\tilde{T}_1) = N \left(\frac{1}{\lambda} + \bar{\tau} \right) \quad (2.8)$$

¹⁰A *cumulative process* is a sequence of partial sums of another process.

¹¹A *Markov renewal-reward process* uses a Markov renewal process to extend the indexing of a cumulative process from \mathbb{N} to $\mathbb{R}_{\geq 0}$.

It is clearly the case that

$$\mathbb{E} \left(\tilde{G}^N + \tilde{C}^N \right) = N \mathbb{E} \left(\tilde{G}_1 + \tilde{C}_1 \right) = N\bar{g} \quad (2.9)$$

Also, for all $t \in \mathbb{R}_{\geq 0}$,

$$\mathbb{E} \left(\tilde{G}(t) \right) = \mathbb{E} (N(t)) \mathbb{E} \left(\tilde{G}_1 \right) = \mathbb{E} (N(t)) \left(\bar{g} - \bar{c} - \frac{c^s}{\lambda} \right) \quad (2.10)$$

$$\mathbb{E} \left(\tilde{C}(t) \right) = \mathbb{E} (N(t)) \mathbb{E} \left(\tilde{C}_1 \right) = \mathbb{E} (N(t)) \left(\bar{c} + \frac{c^s}{\lambda} \right) \quad (2.11)$$

$$\mathbb{E} \left(\tilde{T}(t) \right) = \mathbb{E} (N(t)) \mathbb{E} \left(\tilde{T}_1 \right) = \mathbb{E} (N(t)) \left(\frac{1}{\lambda} + \bar{\tau} \right) \quad (2.12)$$

and, clearly,

$$\mathbb{E} \left(\tilde{G}(t) + \tilde{C}(t) \right) = \mathbb{E} (N(t)) \mathbb{E} \left(\tilde{G}_1 + \tilde{C}_1 \right) = \mathbb{E} (N(t)) \bar{g} \quad (2.13)$$

Stochastic Limits of Net Gain Processes

It can be shown¹² that there exists an $N \in \mathbb{N}$ such that $\mathbb{E}(1/\tilde{T}^N) < \infty$, and so by results of Johns and Miller [28],

$$\text{aslim}_{t \rightarrow \infty} \frac{\tilde{G}(t)}{t} = \lim_{t \rightarrow \infty} \frac{\mathbb{E} \left(\tilde{G}(t) \right)}{t} = \text{aslim}_{N \rightarrow \infty} \frac{\tilde{G}^N}{\tilde{T}^N} = \lim_{N \rightarrow \infty} \mathbb{E} \left(\frac{\tilde{G}^N}{\tilde{T}^N} \right) = \frac{\mathbb{E} \left(\tilde{G}_1 \right)}{\mathbb{E} \left(\tilde{T}_1 \right)} \quad (2.14)$$

This result is frequently used in classical OFT. The ratio $\mathbb{E}(G_1)/\mathbb{E}(T_1)$ may be called the *long-term (average) rate of net gain* and is expressed by

$$\frac{\mathbb{E} \left(\tilde{G}_1 \right)}{\mathbb{E} \left(\tilde{T}_1 \right)} = \frac{\bar{g} - \bar{c} - \frac{c^s}{\lambda}}{\frac{1}{\lambda} + \bar{\tau}} = \frac{\lambda(\bar{g} - \bar{c}) - c^s}{1 + \lambda\bar{\tau}} = \frac{\sum_{i=1}^n \lambda_i p_i (g_i - c_i \tau_i) - c^s}{1 + \sum_{i=1}^n \lambda_i p_i \tau_i}$$

So,

$$\frac{\mathbb{E} \left(\tilde{G}_1 \right)}{\mathbb{E} \left(\tilde{T}_1 \right)} = \frac{\mathbb{E} \left(\tilde{G}^N \right)}{\mathbb{E} \left(\tilde{T}^N \right)} = \frac{\mathbb{E} \left(\tilde{G}(t) \right)}{\mathbb{E} \left(\tilde{T}(t) \right)} \quad (2.15)$$

for all $N \in \mathbb{N}$ and $t \in \mathbb{R}_{>0}$.

¹²See [Appendix A](#).

Variance Under Pseudo-Deterministic Conditions

The statistics of the processes (\tilde{G}^N) , (\tilde{C}^N) , (\tilde{T}^N) , and $(\tilde{G}^N + \tilde{C}^N)$ are of particular interest to us. The expectation of the random variables in these processes are given in Equations (2.6), (2.7), (2.8), and (2.9), respectively; however, it is useful to know their variances¹³ as well, especially when considering risk. Because these four processes are collections of i.i.d. random variables,

$$\begin{aligned}\text{var}(\tilde{G}^N) &= N \text{var}(\tilde{G}_1) = N (\text{var}(\vartheta_1 - \mathcal{U}_1) + \text{var}(C^s \Upsilon_1)) \\ \text{var}(\tilde{C}^N) &= N \text{var}(\tilde{C}_1) = N (\text{var}(\mathcal{U}_1) + \text{var}(C^s \Upsilon_1)) \\ \text{var}(\tilde{T}^N) &= N \text{var}(\tilde{T}_1) = N (\text{var}(\Upsilon_1) + \text{var}(C^s \Upsilon_1)) \\ \text{var}(\tilde{G}^N + \tilde{C}^N) &= N \text{var}(\tilde{G}_1 + \tilde{C}_1) = N \text{var}(\vartheta_1)\end{aligned}$$

for all $N \in \mathbb{N}$. However, the derivations of the variances of \tilde{G}_1 , \tilde{C}_1 , \tilde{T}_1 , and $\tilde{G}_1 + \tilde{C}_1$ are difficult in general. Additionally, they require us to introduce parameters representing the variance of the random variables g_1^i , c_1^i and τ_1^i for all $i \in \{1, 2, \dots, n\}$, which may not be known in applications. Thus, we focus on one particular simplified case; for all $i \in \{1, 2, \dots, n\}$, we assume that

$$\Pr(g_1^i = g_i) = \Pr(c_1^i = c_i) = \Pr(\tau_1^i = \tau_i) = 1$$

This roughly means that the gains, cost rates, and processing times for tasks of any particular type are all deterministic. We also make use of the following assumptions.

- (i) For all $i \in \{1, 2, \dots, n\}$, X_1^i is uncorrelated with each of $(g_1^i - c_1^i \tau_1^i)^2$, $(c_1^i \tau_1^i)^2$, and $(\tau_1^i)^2$.

¹³For a random variable X , the variance $\text{var}(X)$ is $E((X - E(X))^2)$, which is equivalent to $E(X^2) - E(X)^2$. Variance is sometimes called the *second central moment* because it integrates the squared differences from the mean (i.e., the center of the distribution). This is a measure of the likely variability of outcomes.

(ii) For all $i \in \{1, 2, \dots, n\}$, g_1^i is uncorrelated with $c_1^i \tau_1^i$.

(iii) $\varrho_1 - \mathcal{U}_1$ is uncorrelated with $C^s \Upsilon_1$.

(iv) $(C^s)^2$ is uncorrelated with $(\Upsilon_1)^2$.

(v) $(C^s \Upsilon_1)^2$ is independent of I_1 .

From these assumptions, we derive the second moments

$$\mathbb{E}(g^2) = \sum_{i=1}^n \frac{\lambda_i}{\lambda} p_i (g_i)^2 \quad (2.16)$$

$$\mathbb{E}(c^2) = \sum_{i=1}^n \frac{\lambda_i}{\lambda} p_i (c_i \tau_i)^2 \quad (2.17)$$

$$\mathbb{E}(\tau^2) = \sum_{i=1}^n \frac{\lambda_i}{\lambda} p_i (\tau_i)^2 \quad (2.18)$$

$$\mathbb{E}((g - c)^2) = \sum_{i=1}^n \frac{\lambda_i}{\lambda} p_i (g_i - c_i \tau_i)^2 \quad (2.19)$$

which can be used to derive other second moments and variances. So, for all $N \in \mathbb{N}$,

$$\mathbb{E}(\tilde{G}_1^2) = \mathbb{E}((g - c)^2) - 2 \frac{c^s}{\lambda} \mathbb{E}(\tilde{G}_1) \quad (2.20)$$

$$\mathbb{E}(\tilde{C}_1^2) = \mathbb{E}(c^2) - 2 \frac{c^s}{\lambda} \mathbb{E}(\tilde{C}_1) \quad (2.21)$$

$$\mathbb{E}(\tilde{T}_1^2) = \mathbb{E}(\tau^2) - 2 \frac{1}{\lambda} \mathbb{E}(\tilde{T}_1) \quad (2.22)$$

$$\mathbb{E}\left(\left(\tilde{G}_1 + \tilde{C}_1\right)^2\right) = \mathbb{E}(g^2) \quad (2.23)$$

and

$$\text{var}(\tilde{G}^N) = N \left(\text{var}(g - c) + \left(\frac{c^s}{\lambda}\right)^2 \right) \quad (2.24)$$

$$\text{var}(\tilde{C}^N) = N \left(\text{var}(c) + \left(\frac{c^s}{\lambda}\right)^2 \right) \quad (2.25)$$

$$\text{var}(\tilde{T}^N) = N \left(\text{var}(\tau) + \left(\frac{1}{\lambda}\right)^2 \right) \quad (2.26)$$

$$\text{var}(\tilde{G}^N + \tilde{C}^N) = N \text{var}(g) \quad (2.27)$$

Under these assumptions, the only variance in the model comes from the varying time spent searching for tasks and the uncertainty in the type of task encountered.

2.3 Finite Lifetime Analysis: Processing-Based Approach

Recall that the agent suffers no recognition cost upon an encounter with a task. Therefore, it makes sense to exclude tasks that are ignored (i.e., not chosen for processing) from the model entirely by adjusting the encounter rate for each task type. This adjustment is possible in our model specifically because encounters are generated by Poisson processes. Thus, in our approach, we split the task-type processes immediately to thin them of their ignored tasks. We then merge these n thinned processes to form a merged process generated by only the task encounters that result in processing. We can then proceed in the same way as the classical OFT approach, except we assume the agent processes every task from this merged process. Thus, we call this a *split before merge* approach. This approach differs from the classical OFT approach which splits based on processing *after* merging the task-type processes. Because the approach proceeds in an identical way as classical OFT after these modifications, most of this section provides results without a great deal of justification.

2.3.1 Poisson Encounters of Processed Tasks of One Type

For all $i \in \{1, 2, \dots, n\}$, define $(M_i^p(t_s) : t_s \in \mathbb{R}_{\geq 0})$ and $\lambda_i^p \in \mathbb{R}_{>0}$,

$$M_i^p(t_s) \triangleq \sum_{i=1}^{M_i(t_s)} X_i \quad \text{and} \quad \lambda_i^p \triangleq p_i \lambda_i$$

for all $t_s \in \mathbb{R}_{\geq 0}$ and $\zeta \in \mathcal{U}$. Also define \mathcal{G}^p with $\mathcal{G}^p \triangleq \{i \in \{1, 2, \dots, n\} : p_i > 0\}$.

Roughly speaking, for all $\zeta \in \mathcal{U}$, $M_i^p(t_s)$ is a version of $M_i(t_s)$ with all task encounters that do not result in processing removed; that is, $M_i^p(t_s)$ is the number of tasks

processed after searching for t_s time. For all $i \in \mathcal{G}^p$, $(M_i^p(t_s) : t_s \in \mathbb{R}_{\geq 0})$ is a split Poisson process with rate λ_i^p . Therefore, for all $i \in \mathcal{G}^p$, define (\hat{G}_M^i) , (\hat{C}_M^i) , (\hat{T}_M^i) , and (\hat{I}_M^i) as thinned versions of (G_M^i) , (C_M^i) , (T_M^i) , and (I_M^i) respectively. For all $i \in \{1, 2, \dots, n\}$ with $i \notin \mathcal{G}^p$, define $\hat{G}_M^i = \hat{C}_M^i = \hat{T}_M^i = 0$ and $\hat{I}_M^i = i$ for all $M \in \mathbb{N}$. Now we may proceed in an identical way as classical OFT using these thinned processes; however, because the p_i parameter has been absorbed into λ_i^p , it can be omitted.

Poisson Encounters of All Processed Tasks

Assume that $\mathcal{G}^p \neq \emptyset$. This assumption follows from the requirement that an agent must process some finite number of tasks in its lifetime. Define $(M^p(t_s) : t_s \in \mathbb{R}_{\geq 0})$ and $\lambda^p \in \mathbb{R}_{>0}$ with

$$M^p(t_s) \triangleq \sum_{i \in \mathcal{G}^p} M_i^p(t_s) = \sum_{i=1}^n M_i^p(t_s) \quad \text{and} \quad \lambda^p \triangleq \sum_{i \in \mathcal{G}^p} \lambda_i^p = \sum_{i=1}^n \lambda_i^p$$

for all $t_s \in \mathbb{R}_{\geq 0}$ and all $\zeta \in \mathcal{U}$. $(M^p(t_s) : t_s \in \mathbb{R}_{\geq 0})$ is a merged Poisson process with rate λ^p . The process is generated only by encounters that lead to processing. That is, for all $\zeta \in \mathcal{U}$, $M^p(t_s)$ is the *total* number of tasks *processed* after *searching* for t_s time. Call the interevent time process for this task (Υ_m^p) . Therefore, $E(\Upsilon_1^p) = 1/\lambda^p$, $\text{aslim}_{t_s \rightarrow \infty} M^p(t_s) = \infty$, and $E(M^p(t_s)) = \lambda^p t_s$ for all $t_s \in \mathbb{R}_{\geq 0}$.

Merged Task-Type Processes

Define the random processes (\mathcal{D}_M^p) , (\mathcal{U}_M^p) , (\mathcal{T}_M^p) , and (I_M^p) as merged versions of the families $((\hat{G}_M^i)_{i=1}^n)$, $((\hat{C}_M^i)_{i=1}^n)$, $((\hat{T}_M^i)_{i=1}^n)$, and $((\hat{I}_M^i)_{i=1}^n)$ respectively. Each of these processes is an i.i.d. sequence of random variables, where I_1^p and Υ_1^p are assumed to be independent. We use the notations g^p , c^p , and τ^p defined by

$$g^p \triangleq \mathcal{D}_1^p \quad \text{and} \quad c^p \triangleq \mathcal{U}_1^p \quad \text{and} \quad \tau^p \triangleq \mathcal{T}_1^p$$

These respectively represent the gain, cost, and time from processing during a single processing renewal cycle. We also define the symbols \bar{g}^p , \bar{c}^p , and $\bar{\tau}^p$ with

$$\bar{g}^p \triangleq E(g^p) = E(\mathcal{D}_1^p) \quad \text{and} \quad \bar{c}^p \triangleq E(c^p) = E(\mathcal{U}_1^p) \quad \text{and} \quad \bar{\tau}^p \triangleq E(\tau^p) = E(\mathcal{T}_1^p)$$

respectively. Therefore,

$$\bar{g}^p = \sum_{i=1}^n \frac{\lambda_i^p}{\lambda^p} g_j \quad \text{and} \quad \bar{c}^p = \sum_{i=1}^n \frac{\lambda_i^p}{\lambda^p} c_j \tau_j \quad \text{and} \quad \bar{\tau}^p = \sum_{i=1}^n \frac{\lambda_i^p}{\lambda^p} \tau_j$$

So, these expectations are weighted sums of parameters. In particular, if $n = 1$ (and $p_1 = 1$),

$$\bar{g}^p = g_1(\tau_1) \quad \text{and} \quad \bar{c}^p = c_1 \tau_1 \quad \text{and} \quad \bar{\tau}^p = \tau_1$$

This result is useful when visualizing optimization results. Additionally,

$$E(C^s \Upsilon_1^p | I_1^p = i) = E(C^s \Upsilon_1^p) = \frac{c^s}{\lambda^p}$$

We will use these results frequently in expressions of statistics of interest.

2.3.2 Process-Only Markov Renewal Process

Define i.i.d. random processes (G_{N^p}) , (C_{N^p}) , and (T_{N^p}) with

$$G_{N^p} \triangleq \mathcal{D}_{N^p}^p - \mathcal{U}_{N^p}^p - C^s \Upsilon_{N^p}^p$$

$$C_{N^p} \triangleq \mathcal{U}_{N^p}^p + C^s \Upsilon_{N^p}^p$$

$$T_{N^p} \triangleq \mathcal{T}_{N^p}^p + \Upsilon_{N^p}^p$$




Figure 2.2: The process-only Markov renewal process, where the solid dot is the renewal point that starts each cycle.

for all $N^p \in \mathbb{N}$ and $\zeta \in \mathcal{U}$. Clearly, the i.i.d. process $(G_{N^p} + C_{N^p})$ has $G_{N^p} + C_{N^p} = \mathfrak{D}_{N^p}^p$

for all $N^p \in \mathbb{N}$ and $\zeta \in \mathcal{U}$. Also, $\Pr(T_1 = 0) = 0$ and

$$\mathbb{E}(G_1) = \sum_{i=1}^n \frac{\lambda_i^p}{\lambda^p} (g_i - c_i \tau_i) - \frac{c^s}{\lambda^p} \quad (2.28)$$

$$\mathbb{E}(C_1) = \sum_{i=1}^n \frac{\lambda_i^p}{\lambda^p} c_i \tau_i + \frac{c^s}{\lambda^p} \quad (2.29)$$

$$\mathbb{E}(T_1) = \frac{1}{\lambda^p} + \sum_{i=1}^n \frac{\lambda_i^p}{\lambda^p} \tau_i \quad (2.30)$$

$$\mathbb{E}(G_1 + C_1) = \sum_{i=1}^n \frac{\lambda_i^p}{\lambda^p} g_i \quad (2.31)$$

Because $0 < \mathbb{E}(T_1) < \infty$, define $(N^p(t) : t \in \mathbb{R}_{\geq 0})$ with

$$N^p(t) \triangleq \sup \left\{ N^p \in \mathbb{N} : \sum_{i=1}^{N^p} T_i \leq t \right\}$$

for all $t \in \mathbb{R}_{\geq 0}$ and $\zeta \in \mathcal{U}$. That is, for all $\zeta \in \mathcal{U}$, $N^p(t)$ is the *total* number of tasks *processed* from time 0 to *total* time t . This is a Markov renewal process depicted in [Figure 2.2](#), and one iteration around this process will be known as a *processing cycle*. The holding time for this process *always* includes *both* search and processing time. Clearly, (T_{N^p}) is the interevent time process, as $\lim_{N^p \rightarrow \infty} N^p(t) = \infty$, and $0 \leq \mathbb{E}(N^p(t)) \leq \lambda^p t$ for all $t \in \mathbb{R}_{\geq 0}$.

Cumulative Reward Processes and Their Statistics

Define the cumulative processes (G^{N^p}) , (C^{N^p}) , and (T^{N^p}) with

$$G^{N^p} \triangleq \sum_{i=1}^{N^p} G_i \quad \text{and} \quad C^{N^p} \triangleq \sum_{i=1}^{N^p} C_i \quad \text{and} \quad T^{N^p} \triangleq \sum_{i=1}^{N^p} T_i$$

and the Markov renewal-reward processes $(G(t) : t \in \mathbb{R}_{\geq 0})$, $(C(t) : t \in \mathbb{R}_{\geq 0})$, and $(T(t) : t \in \mathbb{R}_{\geq 0})$ with

$$G(t) \triangleq G^{N^p(t)} \quad \text{and} \quad C(t) \triangleq C^{N^p(t)} \quad \text{and} \quad T(t) \triangleq T^{N^p(t)}$$

Clearly, processes $(G^{N^p} + C^{N^p})$ and $(G(t) + C(t) : t \in \mathbb{R}_{\geq 0})$ are well-defined. Therefore, for all $N^p \in \mathbb{N}$

$$\mathbb{E}(G^{N^p}) = N^p \mathbb{E}(G_1) \quad \text{and} \quad \mathbb{E}(C^{N^p}) = N^p \mathbb{E}(C_1) \quad \text{and} \quad \mathbb{E}(T^{N^p}) = N^p \mathbb{E}(T_1)$$

and so $\mathbb{E}(G^{N^p} + C^{N^p}) = N^p \mathbb{E}(G_1 + C_1)$. Also, for all $t \in \mathbb{R}_{\geq 0}$,

$$\mathbb{E}(G(t)) = \mathbb{E}(N^p(t)) \mathbb{E}(G_1)$$

$$\mathbb{E}(C(t)) = \mathbb{E}(N^p(t)) \mathbb{E}(C_1)$$

$$\mathbb{E}(T(t)) = \mathbb{E}(N^p(t)) \mathbb{E}(T_1)$$

and so $\mathbb{E}(G(t) + C(t)) = \mathbb{E}(N^p(t)) \mathbb{E}(G_1 + C_1)$.

Limits of Cumulative Reward Processes

There exists¹⁴ an $N^p \in \mathbb{N}$ such that $\mathbb{E}(1/N^p) < \infty$, so

$$\text{aslim}_{t \rightarrow \infty} \frac{G(t)}{t} = \lim_{t \rightarrow \infty} \frac{\mathbb{E}(G(t))}{t} = \text{aslim}_{N^p \rightarrow \infty} \frac{G^{N^p}}{T^{N^p}} = \lim_{N^p \rightarrow \infty} \mathbb{E} \left(\frac{G^{N^p}}{T^{N^p}} \right) = \frac{\mathbb{E}(G_1)}{\mathbb{E}(T_1)} \quad (2.32)$$

and

$$\text{aslim}_{t \rightarrow \infty} \frac{N^p(t)}{t} = \lim_{t \rightarrow \infty} \frac{\mathbb{E}(N^p(t))}{t} = \text{aslim}_{N^p \rightarrow \infty} \frac{N^p}{T^{N^p}} = \lim_{N^p \rightarrow \infty} \mathbb{E} \left(\frac{N^p}{T^{N^p}} \right) = \frac{1}{\mathbb{E}(T_1)} \quad (2.33)$$

¹⁴See [Appendix A](#).

The ratio $E(G_1)/E(T_1)$ may be called the *long-term (average) rate of net gain* and has the expression

$$\frac{E(G_1)}{E(T_1)} = \frac{\overline{g^p} - \overline{c^p} - \frac{c^s}{\lambda^p}}{\frac{1}{\lambda^p} + \overline{\tau^p}} = \frac{\sum_{i=1}^n \lambda_i^p (g_i - c_i \tau_i) - c^s}{1 + \sum_{i=1}^n \lambda_i^p \tau_i} = \frac{\lambda^p (\overline{g^p} - \overline{c^p}) - c^s}{1 + \lambda^p \overline{\tau^p}}$$

So,

$$\frac{E(G_1)}{E(T_1)} = \frac{E(G^N)}{E(T^N)} = \frac{E(G(t))}{E(T(t))} \quad (2.34)$$

for all $N \in \mathbb{N}$ and $t \in \mathbb{R}_{>0}$. Additionally, $E(G_1)/E(T_1) = E(\tilde{G}_1)/E(\tilde{T}_1)$, which shows an important connection between this approach and the classical OFT approach.

Variance Under Pseudo-Deterministic Conditions

To define the variance of (G^{N^p}) , (C^{N^p}) , (T^{N^p}) , and $(G^{N^p} + C^{N^p})$, we must again assume that $\Pr(g_1^i = g_i) = \Pr(c_1^i = c_i) = \Pr(\tau_1^i = \tau_i) = 1$ and that

- (i) For all $i \in \{1, 2, \dots, n\}$, X_1^p is uncorrelated with each of $(g_1^i - c_1^i \tau_1^i)^2$, $(c_1^i \tau_1^i)^2$, and $(\tau_1^i)^2$.
- (ii) \mathfrak{D}_1^p is uncorrelated with $C^s \Upsilon_1^p$.
- (iii) $(C^s \Upsilon_1^p)^2$ is independent of I_1^p .
- (iv) $(C^s)^2$ is uncorrelated with $(\Upsilon_1^p)^2$.
- (v) For all $i \in \{1, 2, \dots, n\}$, g_1^i is uncorrelated with $c_1^i \tau_1^i$.

These assumptions yield the second moments

$$\mathbb{E}((g^p)^2) = \sum_{i=1}^n \frac{\lambda_i^p}{\lambda^p} (g_i)^2 \quad (2.35)$$

$$\mathbb{E}((c^p)^2) = \sum_{i=1}^n \frac{\lambda_i^p}{\lambda^p} (c_i \tau_i)^2 \quad (2.36)$$

$$\mathbb{E}((\tau^p)^2) = \sum_{i=1}^n \frac{\lambda_i^p}{\lambda^p} (\tau_i)^2 \quad (2.37)$$

$$\mathbb{E}((g^p - c^p)^2) = \sum_{i=1}^n \frac{\lambda_i^p}{\lambda^p} (g_i - c_i \tau_i)^2 \quad (2.38)$$

which can be used to derive variances and other second moments. In particular, for all $N^p \in \mathbb{N}$,

$$\mathbb{E}(G_1^2) = \mathbb{E}((g^p - c^p)^2) - 2 \frac{c^s}{\lambda^p} \mathbb{E}(G_1) \quad (2.39)$$

$$\mathbb{E}(C_1^2) = \mathbb{E}((c^p)^2) - 2 \frac{c^s}{\lambda^p} \mathbb{E}(C_1) \quad (2.40)$$

$$\mathbb{E}(T_1^2) = \mathbb{E}((\tau^p)^2) - 2 \frac{1}{\lambda^p} \mathbb{E}(T_1) \quad (2.41)$$

$$\mathbb{E}((G_1 + C_1)^2) = \mathbb{E}((g^p)^2) \quad (2.42)$$

and

$$\text{var}(G^{N^p}) = N^p \left(\text{var}(g^p - c^p) + \left(\frac{c^s}{\lambda^p} \right)^2 \right) \quad (2.43)$$

$$\text{var}(C^{N^p}) = N^p \left(\text{var}(c^p) + \left(\frac{c^s}{\lambda^p} \right)^2 \right) \quad (2.44)$$

$$\text{var}(T^{N^p}) = N^p \left(\text{var}(\tau^p) + \left(\frac{1}{\lambda^p} \right)^2 \right) \quad (2.45)$$

$$\text{var}(G^{N^p} + C^{N^p}) = N^p \text{var}(g^p) \quad (2.46)$$

2.4 Relationship Between Analysis Approaches

Recall that for all $i \in \{1, \dots, n\}$, $\lambda_i^p = p_i \lambda_i$. Keeping this in mind, it is clear that in general (i.e., for any $t \in \mathbb{R}_{\geq 0}$, $N, N^p \in \mathbb{W}$)

$$\mathbb{E}(G(t)) \neq \mathbb{E}(G^{N^p}) \neq \mathbb{E}(G_1) \neq \mathbb{E}(\tilde{G}_1) \neq \mathbb{E}(\tilde{G}^N) \neq \mathbb{E}(\tilde{G}(t))$$

and

$$\mathbb{E}(T(t)) \neq \mathbb{E}(T^{N^p}) \neq \mathbb{E}(T_1) \neq \mathbb{E}(\tilde{T}_1) \neq \mathbb{E}(\tilde{T}^N) \neq \mathbb{E}(\tilde{T}(t))$$

However,

$$\frac{\mathbb{E}(G(t))}{\mathbb{E}(T(t))} = \frac{\mathbb{E}(G^{N^p})}{\mathbb{E}(T^{N^p})} = \frac{\mathbb{E}(G_1)}{\mathbb{E}(T_1)} = \frac{\mathbb{E}(\tilde{G}_1)}{\mathbb{E}(\tilde{T}_1)} = \frac{\mathbb{E}(\tilde{G}^N)}{\mathbb{E}(\tilde{T}^N)} = \frac{\mathbb{E}(\tilde{G}(t))}{\mathbb{E}(\tilde{T}(t))} \quad (2.47)$$

for all $t \in \mathbb{R}_{> 0}$ and $N, N^p \in \mathbb{N}$. Note the following.

- (i) $\mathbb{E}(\tilde{T}_1) > 0$ and $\mathbb{E}(T_1) > 0$, and so all of the ratios in [Equations \(2.47\)](#) are well-defined.
- (ii) There are no restrictions on the sign of $\mathbb{E}(\tilde{G}_1)$ or $\mathbb{E}(G_1)$. These can be negative, zero, or positive.
- (iii) There are no restrictions on the sign of $\mathbb{E}(\tilde{C}_1)$ or $\mathbb{E}(C_1)$. These can be negative, zero, or positive.

Points (ii) and (iii) allow for flexible interpretations of *gain* and *cost*. With the appropriate assignment of signs, gains can be viewed as time-invariant costs, and costs can be viewed as time-varying gains. This shows the flexibility of this generalized model.

The equalities in Equation (2.47) imply that the stochastic limits in Equation (2.5) are equal to the stochastic limits in Equation (2.33); regardless of approach, the long-term rate of net point gain is equivalent. For any number of processing cycles or OFT cycles completed, the ratio of expected net gain to expected time will be equal. Processing is guaranteed in a processing cycle, so a single processing cycle has a higher expected net gain than a single OFT cycle; however, the expected holding time of a processing cycle is longer because encounters with ignored tasks are included as part of the cycle's holding time. Thus, the ratio of expected net gain to expected time is the same for cycles of either type.

2.5 Weaknesses of the Model

Several features are not included in the model.

Rates and Costs: Recognition costs, variable search rates, and variable processing rates are not modeled. Also, although encounters are assumed to happen at random, they are assumed to be driven by a homogenous Poisson process (i.e., the average rate of encounters is time-invariant).

Perfect Detection: When an agent encounters a task, its behavior depends upon the type of that task. The model assumes that the agent can detect task types with no error. This model has been built so that it may potentially be augmented with support for detection error.

Linear Cost Model: All costs are assumed to be linear in time in this model. Thus, given any interval of time, the cost of that interval of time is assumed to be the product of the length of that interval with some constant, which we call a *cost*

rate. In most cases, that rate need not be deterministic; however, it must be uncorrelated with the interval of time.

Known Search Cost Rate: Search costs are also assumed by linear with respect to time; however, they are also assumed to be deterministic. This assumption is necessary to use the results from renewal theory that are central to classical OFT methods. Thus, in many cases where these results are not used, this deterministic assumption can be relaxed.

Competition and Cooperation: The direct effect of other agents (e.g., competition or cooperation) on the environment is not modeled here in any specific way. Cody [19] views this as a weakness of the early solitary foraging models and introduces an optimal diet model that incorporates multiple foragers competing for resources. However, the parameters of the [Cody](#) model are too abstract to be specified with physical quantities, and each forager in the model has a coarse set of behavioral options. Additionally, many engineering applications fit the solitary model well (e.g., autonomous surveillance vehicles).

State Dependency: Our model is not state-dependent. That is, the reaction of an agent to an encounter does not change over its lifetime (i.e., it is a static model). Schoener [55] documents many cases where foragers adjust their behavior when satiated. Houston and McNamara [24] handle state-dependent behaviors mathematically and show that they will often be advantageous when compared to static behaviors. However, in engineering applications it may be desirable to have behaviors that do not change over time. For example, if the computational abilities of an agent are limited, complex state-dependent behavior may

not be possible. There may also be biological examples where dynamic adaptations based on feedback are not feasible. Thus, optimization over a set of time-invariant behaviors may be desirable in a number of applications.

Despite the limitations of the model, it is sufficiently generic to have utility in a wide range of applications. Adding any further complexity to the model may make solutions too complex to be practical for implementation.

CHAPTER 3

STATISTICAL OPTIMIZATION OBJECTIVES FOR SOLITARY BEHAVIOR

The efficacy of any particular behavior may be measured quantitatively in various ways. In this chapter, we approach the problem of combining appropriate statistics so that the utility of solitary behaviors can be measured for a given application. Choosing a static behavior to maximize some unit of expected value is analogous to choosing investments to maximize future returns. Reflecting this analogy, behavioral ecology has borrowed methods from investment theory and capital budgeting for behavioral analysis. We also use these methods, collectively known as modern portfolio theory (MPT), to analyze our model; however, we generalize the classical OFT approach. This approach not only allows it to be applied to engineering problems, but it also provides answers to some of the criticisms of the theory. Additionally, we suggest new ways of describing optimal agent behavior and relationships among existing methods.

The major purpose of this chapter is to introduce functions that combine statistics of the agent model to measure the utility of solitary behaviors. Behaviors that maximize these functions may be called *optimal*. In [Section 3.1](#), we define the structure of the optimization functions that are interesting to us. In [Section 3.2](#), we describe the

optimization approach used frequently in classical OFT. In [Section 3.3](#), we propose an alternate approach and give new or refined optimization objectives for analyzing agent behavior. Finally, in [Section 3.4](#), we briefly discuss how insights from post-modern portfolio theory (PMPT) may inspire new optimization approaches in both agent design in engineering and agent analysis in biology. All results discussed in this chapter will be qualitative and justified graphically. Specific analytical optimization results for some of the objectives discussed here are given in [Chapter 4](#).

3.1 Objective Function Structure

Optimization functions usually combine multiple optimization *objectives* in a way that captures the relative value of each of those objectives. In our case, each of our objectives is a statistic taken from the model in [Chapter 2](#). Therefore, in [Section 3.1.1](#), we present statistics that could serve as objectives for optimization and methods for combining them. In [Section 3.1.2](#), we discuss motivations for constraining the set of feasible behaviors and show how these constrained sets can be incorporated into optimization. Finally, in [Section 3.1.3](#), we discuss the importance of exploring a variety of optimization criteria.

3.1.1 Statistics of Interest

[Table 3.1](#) shows some obvious choices for statistics to be used as optimization objectives. However, other statistics like $E(G^N/T^N)$ (i.e., average gain per unit time) or $E((G^N + C^N)/C^N)$ (i.e., average efficiency) for all $N \in \mathbb{N}$ could also be relevant. Economists [e.g., [17](#), [29](#), [30](#), [31](#), [63](#)] might argue that the skewness¹ of each of these

¹For a random variable X , its *skewness* is a measure of the symmetry of its (Lebesgue) probability density f_X . The standard definition of skewness is $E((X - E(X))^3)/\text{std}(X)^3$. Note that this is a scaled version of the third central moment.

	Means		Variances	
Net Gain Statistics:	$E(G_1)$	$E(\tilde{G}_1)$	$\text{var}(G_1)$	$\text{var}(\tilde{G}_1)$
Cost Statistics:	$E(C_1)$	$E(\tilde{C}_1)$	$\text{var}(C_1)$	$\text{var}(\tilde{C}_1)$
Time Statistics:	$E(T_1)$	$E(\tilde{T}_1)$	$\text{var}(T_1)$	$\text{var}(\tilde{T}_1)$

Table 3.1: Common statistics used in optimization of solitary agent behavior.

random variables would be a reasonable statistic to study because it may be desirable to have random variables that are distributed asymmetrically (e.g., net gains that are more often high than low)². Of course, any one of these statistics may not capture all relevant objectives of a problem. For example, it may be desirable to maximize both $E(G_1)$ and $-E(T_1)$ (i.e., minimize $E(T_1)$); however, it may not be possible to accomplish both of these simultaneously. Therefore, here we discuss the construction of compound objectives that allow for optimization with respect to multiple criteria.

Take a problem with $m \in \mathbb{N}$ relevant optimization objectives. For all objective functions to be minimized, replace the function with its additive or multiplicative inverse (i.e., replace a function f with the function $-f$ or, for functions with strictly positive or strictly negative ranges, $1/f$); therefore, the ideal objective is to maximize all m functions. Collect these m objective functions into m -vector x where $x = \{x_1, x_2, \dots, x_m\}$. Use the weighting vector $w \in \mathbb{R}_{\geq 0}^m$ with $w = \{w_1, w_2, \dots, w_m\}$ to represent the relative value of each of these objectives. Therefore, the compound objective functions

$$w_1x_1 + w_2x_2 + \dots + w_mx_m \quad \text{or} \quad \min\{w_1x_1, w_2x_2, \dots, w_mx_m\} \quad (3.1)$$

²This might be called skewness preference. It is also desirable to optimize skewness simply to prevent deleterious asymmetry.

represent different ways to combine all m objectives. The former of these two compound objectives is a linear combination of statistics (i.e., $w^\top x$), and an optimal behavior for this function will be Pareto efficient³ with respect to the m objective functions. Maximization of the latter of these two compound objectives represents a *maximin* optimization problem. Lagrange multiplier methods (i.e., Karush-Khunan-Tucker (KKT) conditions) [10] can be used to study the optimal solutions to both forms in Equation (3.1).

3.1.2 Optimization Constraints

In a given foraging problem, it is not necessarily the case that all modeled behaviors are applicable or even possible. That is, optimization analysis must be considered with respect to a set of feasible behaviors. The following are some examples of constraints that have been found in the literature; suggestions for how those constraints could be implemented in this model are also given.

Time Constraints: The economics-inspired graphical foraging model of Rapport [51] considers level *indifference curves* of an energy function. Each of these curves represents a set of combinations of prey where each combination returns the same energetic gain to the forager. Rapport then assumes that the forager has a finite lifetime and surrounds all prey combinations that can be completed in this time with a boundary called the *consumption frontier*⁴. The optimal

³To be *Pareto efficient* or *Pareto optimal* means that any deviation that yields an increase in one objective function will also result in a decrease in another objective function. Pareto optimal solutions characterize tradeoffs in optimization objectives. If deviation from some behavior will increase all objective functions, then that behavior cannot be Pareto efficient. The set of all Pareto efficient solutions is called the *Pareto frontier*.

⁴The consumption frontier is a Pareto frontier. Diets on this frontier return the greatest gain for their foraging time.

diet combination is the point of tangency between the consumption frontier and some indifference curve. In other words, this is the combination of prey items that returns the highest energetic gain for the given finite lifetime. We can quantify this idea by maximizing $E(G(t))$ subject to the constraint $t \leq \bar{T}$ where $\bar{T} \in \mathbb{R}_{>0}$. Because [Rapport](#) gives a qualitative explanation for the observations in Murdoch [42], the analytical application of our model with this time constraint could give a quantitative explanation.

Nutrient Constraints: Pulliam [48] optimizes a point gain per unit time function similar in form to $E(\tilde{G}_1)/E(\tilde{T}_1)$, but the notion of nutrient constraints is added. That is, there are $m \in \mathbb{N}$ nutrients and all tasks of type $i \in \{1, 2, \dots, n\}$ return quantity ρ_{ij} of nutrient $j \in \{1, 2, \dots, m\}$. Pulliam then calls $M_j \in \mathbb{R}_{\geq 0}$ a minimum amount of nutrient j that must be returned from processing. The goal is to maximize the rate of point gain while maintaining this minimum nutrient level. These nutrient constraints could be added to our model as well. As Pulliam notes, under these constraints, optimal behaviors often include partial preferences. In the unconstrained classical OFT problem, it is sufficient for optimality to either process all or none of tasks of a particular type; however, with nutrient constraints it may be necessary for optimality that only a fraction of the encountered tasks of a certain type be processed⁵.

Encounter-Rate Constraints: Gendron and Staddon [21] and Pavlic and Passino [46]

explore the optimization of a point gain per unit time function as well; however,

⁵In Chapter 4, we generalize the classical OFT result to show that over a closed interval of preference probabilities, sufficiency is associated with the endpoints. The results of Pulliam [48] effectively make that interval a function of nutrition requirements; under these constraints, partial preferences may be necessary for optimality.

the impact of speed choice on imperfect detection is also introduced. That is, with perfect detection, an increase in speed will most likely come with an increase in encounter rate with tasks of every type. However, when detection errors can occur, the relationship between encounter rate and speed may be arbitrarily nonlinear. If this exact relationship is not known, it may be sufficient to restrict search speed to a range where detection is reliable. If the impact of search speed were added to our model (e.g., if encounter-rate was parameterized by speed), this restriction could be modeled as constraints on search speed. The resulting optimal behavior would include a search speed that provides the optimal encounter rates subject to imperfect detection.

Any optimization function of a form in Equation (3.1) subject to a finite number of equality or non-strict inequality constraints⁶ may be analyzed with Lagrange multiplier methods. Therefore, in principle, a wide range of constrained optimization problems can be studied.

3.1.3 Impact of Function Choice on Optimal Behaviors

As discussed in Section 3.2.1, classical OFT results come from maximizing the long-term rate of gain (e.g., $E(\tilde{G}_1)/E(\tilde{T}_1)$). This choice follows from the argument of Pyke et al. [49] that optimizing this long-term rate synthesizes the two extremes, energetic maximization and time minimization, of a general model of foraging given by Schoener [55]. This rate approach is taken by Pulliam [48] whose quantitative results show that the optimal diet predicted by a rate maximizer depend only on the encounter rates with prey types in the diet. However, Rapport [51] focusses only

⁶A *strict* inequality constraint uses $<$ or $>$; therefore, a *non-strict* or *weak* inequality constraint uses \leq and \geq .

on gain maximization (in finite time) and shows that the optimal diet depends on encounter rates with all prey types. These two results are very different, and the only justification for using the first result follows from a purely intuitive argument from Pyke et al. [49]. However, the result from Rapport is entirely valid from a perspective of the foundational work of Schoener. Therefore, it is clear that one optimization criterion will not fit all problems. Clearly, is important to investigate other functions that may be more appropriate for specific problems.

3.2 Classical OFT Approach to Optimization

As discussed by Stephens and Charnov [59], classical OFT approaches optimization from two perspectives which are both based on evolutionary arguments. The first analyzes behaviors that optimize of the asymptotic limit of rate of net gain. The second assumes the agent must meet some energetic requirement and maximizes its probability of success. The former, which we describe in Section 3.2.1, is called *rate maximization*, and the latter, which we describe in Section 3.2.2, is described as being *risk sensitive*. Both approaches develop optimal static behaviors for the solitary agent.

3.2.1 Maximization of Long-Term Rate of Net Gain

In biological contexts, it is expected that natural selection will favor foraging behaviors that provide greater future reproductive success, a common surrogate for Darwinian fitness. So, functions mapping specific behaviors to quantitative measures of reproductive success can be optimized to predict behaviors that should be maintained by natural selection. Schoener [55] defines such a model, and while quantities in the model are too difficult to define for most cases, behaviors predicted by the

model fall on a continuum from foraging time minimizers (when energy is held constant) to energy maximizers (when foraging time is held constant). In other words, behaviors should be excluded if there exists another behavior that has both a higher energy return and a lower time. Pyke et al. [49] argue that the *rate* of net energy intake is the most general function to be maximized as it captures both extremes on the Schoener continuum by asserting an upward pressure on energy intake and a downward pressure on foraging time. This will allow a forager to achieve its energy consumption needs while also leaving it enough time for other activities such as reproduction and predator avoidance. This interpretation is only valid over the space of behaviors with positive net energetic intake. For example, rate maximization puts an upward pressure on foraging time for behaviors that return negative net energetic intake. This is not recognized by Pyke et al., and the continuum of behaviors described by Schoener explicitly exclude these time maximizers. However, from a survival viewpoint, it makes sense that foragers facing a negative energy budget should maximize time foraging. Therefore, rate maximization encapsulates two conditional optimization problems; it trades off net gain and total time in a way that is dependent upon energy reserves.

The rate of net energy intake can be defined in different ways. Using the terms from Chapter 2, it could be defined as $\tilde{G}(t)/t$ or $E(\tilde{G}(t))/t$ for any $t \in \mathbb{R}_{\geq 0}$ or \tilde{G}^N/\tilde{T}^N or $E(\tilde{G}^N/\tilde{T}^N)$ for any $N \in \mathbb{N}$. However, Pyke et al. also argue that rates should be calculated over the entire lifetime of the forager. Thus, rather than taking a particular $t \in \mathbb{R}_{\geq 0}$ or $N \in \mathbb{N}$, the asymptotic limits of these ratios should be taken. Conveniently,

Equation (2.14) shows that all of these limits are equivalent. By Equation (2.15),

$$\begin{aligned} \frac{\mathbb{E}(\tilde{G}_1)}{\mathbb{E}(\tilde{T}_1)} &= \frac{\mathbb{E}(\tilde{G}^{N_*})}{\mathbb{E}(\tilde{T}^{N_*})} = \text{aslim}_{N \rightarrow \infty} \frac{\tilde{G}^N}{\tilde{T}^N} = \lim_{N \rightarrow \infty} \mathbb{E}\left(\frac{\tilde{G}^N}{\tilde{T}^N}\right) \\ &= \frac{\mathbb{E}(\tilde{G}(t_*))}{\mathbb{E}(\tilde{T}(t_*))} = \text{aslim}_{t \rightarrow \infty} \frac{\tilde{G}(t)}{t} = \lim_{t \rightarrow \infty} \frac{\mathbb{E}(\tilde{G}(t))}{t} \end{aligned} \quad (3.2)$$

for any $t_* \in \mathbb{R}_{>0}$ and $N_* \in \mathbb{N}$. For this reason, the ratio of expectations $\mathbb{E}(\tilde{G}_1)/\mathbb{E}(\tilde{T}_1)$ has received significant interest in classical OFT [e.g., 24, 59, 60]. We call this ratio the *long-term (average) rate of net gain*. Note that by Equation (2.47) this ratio plays an identical role in our analysis approach when we consider the asymptotic case.

Opportunity Cost and Pareto Optimality

Houston and McNamara [24] provide an interesting interpretation of $\mathbb{E}(\tilde{G}_1)/\mathbb{E}(\tilde{T}_1)$. They define constant $\tilde{\gamma}^* \in \mathbb{R}$ to be the maximum value of $\mathbb{E}(\tilde{G}_1)/\mathbb{E}(\tilde{T}_1)$ (i.e., the long-term rate of net gain) over the set of feasible agent behaviors. They then treat rate $\tilde{\gamma}^*$ as a factor converting time spent between encounters to maximum points possible from that time. Therefore, $\tilde{\gamma}^*$ converts time into its equivalent *opportunity cost* (i.e., gain paid per unit time). They show that the behavior that maximizes

$$\mathbb{E}(\tilde{G}_1 - \tilde{\gamma}^* \tilde{T}_1) \quad (3.3)$$

will also be the behavior that achieves the maximum long-term rate of gain $\tilde{\gamma}^*$. So, maximizing the long-term rate of gain is equivalent to maximizing the per-cycle *gain* after being discounted by the opportunity cost of the cycle time⁷. Solving for this

⁷There is a related result by Engen and Stenseth [20] that predicts the optimal behavior on simultaneous encounters. This is described by both Houston and McNamara [24] and Stephens and Krebs [60], and Houston and McNamara show this simultaneous encounter result to follow from the opportunity cost result.

behavior can be done analytically only if $\tilde{\gamma}^*$ is known, and so the method of [Houston and McNamara](#) numerically solves for the optimal behavior using iteration, which could be a weakness of this approach. However, it demonstrates an important interpretation of $E(\tilde{G}_1)/E(\tilde{T}_1)$ as the opportunity cost of time. Not surprisingly, this also shows that the behavior that maximizes the long-term rate of gain is Pareto optimal with respect to maximization of $E(\tilde{G}_1)$ and (maximization) minimization of $E(\tilde{T}_1)$ when $\tilde{\gamma}^* > 0$ ($\tilde{\gamma}^* < 0$); that is, this optimal behavior represents a particular tradeoff between net gain and total time. This Pareto interpretation casts $\tilde{\gamma}^*$ as the relative importance of minimizing time, which is consistent with notion of opportunity cost⁸. The numerical approach to finding $\tilde{\gamma}^*$ and the corresponding optimal behavior is equivalent to sliding along a continuum of Pareto efficient solutions (i.e., tradeoffs of net gain and total time).

Equilibrium Renewal Process as an Attractive Alternative

Charnov and Orians [16] note that it is desirable to derive the *equilibrium renewal process* rate of net gain. That is, introduce a $T_1 \in \mathbb{R}_{>0}$ and redefine the process to start after T_1 foraging time has past. Hence, runtime t represents the length of the interval immediately after time T_1 , and so quantity of interest to [Charnov and Orians](#) is $E(G(t))/t$, which represents the average rate of net gain returned to an agent when the agent is in equilibrium with its environment (i.e., after the decay of any initial transients). However, they point out that this rate is only known for such a process if it is additionally assumed that the net gain on each OFT cycle is independent of the total time of each OFT cycle (in particular, the processing time of each cycle). In

⁸When $\tilde{\gamma}^* < 0$, the relative importance of minimizing time is *negative*, which indicates that $|\tilde{\gamma}^*|$ is the relative importance of *maximizing* time (i.e., an opportunity *gain*).

$$\begin{aligned}
\mathbb{E}(\tilde{G}_1) &= g_1(\tau_1) \\
\mathbb{E}(\tilde{T}_1) &= \tau_1 + \frac{1}{\lambda} \\
\tilde{\gamma} &\triangleq \frac{g_1(\tau_1)}{\tau_1 + \frac{1}{\lambda}} \\
\lambda &= \lambda_1 \\
\diamond^* &\triangleq \max \{ \diamond \}
\end{aligned}$$




Figure 3.1: Rate maximization in classical OFT. It is assumed that $n = 1$, $c^s = 0$, and $c_1 = 0$. The constraint that $p_1 = 1$ is also applied. The optimal processing time is denoted t^* , and the corresponding maximal rate is denoted $\tilde{\gamma}^*$ and shown as a slope of a tangent line.

that case, $\mathbb{E}(G(t))/t$ can also be expressed as the ratio $\mathbb{E}(\tilde{G}_1)/\mathbb{E}(\tilde{T}_1)$. Unfortunately, it is rare that net gain and processing time will be independent in a practical system. Analytical results are not available otherwise. For this reason, when $\mathbb{E}(\tilde{G}_1)/\mathbb{E}(\tilde{T}_1)$ is used it is usually assumed to be a limiting case (i.e., a rate over a long time rather than a short-term rate after a long time).

Graphical Interpretation of Rate Maximization

When an agent is only free to choose its (average) processing times, the tasks are said to occur in *patches* or to be *patchily distributed* [60]. Take such a case with a single task type and no search or processing costs (i.e., $n = 1$, $c^s = c_1 = 0$, $p_1 = 1$, and $\tau_1 \in \mathbb{R}_{\geq 0}$). Stephens and Krebs [60] show that this problem has an insightful graphical solution. Consider Figure 3.1. The $g_1(\tau_1)$ function is plotted with respect to feasible choices of τ_1 and a mark is made at the point $(-1/\lambda, 0)$. For any τ_1 , the corresponding long-term rate of gain is the slope of a line that connects points $(-1/\lambda, 0)$ and $(\tau_1, g(\tau_1))$. Therefore, the optimal τ_1 (shown as t^*) is the one that corresponds with the line with the maximal slope, and that slope will be the maximal

long-term rate of gain (shown as $\tilde{\gamma}^*$). In Section 3.3.2, we show how this graphical interpretation can be extended to the general case⁹ (i.e., with multiple types, costly searching and processing, and tasks that may or may not be patchily distributed).

Several conclusions can be drawn from Figure 3.1. For differentiable functions with $g_1(0) = 0$ and $g_1'(0) > 0$, the optimal processing time t^* must be such that $g_1'(t^*)$ is equal to the long-term rate of gain. In particular, if g_1 is a concave function, then this line will be the *unique* tangent line that crosses $(0, 1/\lambda)$. Rate-maximization for the classical OFT model is said to follow the *marginal value theorem (MVT)* [14, 16]. This means that the average time an agent processes patchily distributed tasks of a certain type is the time when the average rate of point gain for the task type drops to the average rate of point gain for the environment. That is, processing should continue until the marginal return from the next instant of processing is less than the environmental average rate of gain¹⁰.

3.2.2 Minimization of Net Gain Shortfall

Because rate maximization depends only on first-order statistics, it disregards the standard deviation¹¹ of random variables in the model. For example, an agent with a behavior that maximizes its long-term rate of net gain may bypass frequently encountered tasks with small gains regardless of any survival needs. However, if the agent must meet a net gain requirement in finite time, it may be beneficial to

⁹We show this interpretation using our approach to defining the relevant statistics of the model; however, our method can also be applied to the classical OFT statistics in an obvious way (i.e., with little more than a change of notation).

¹⁰This interpretation is really only accurate for a deterministic agent model. In the general stochastic agent model, the MVT need only be observed in the first-order statistics of the gains and processing times.

¹¹For random variable X , the standard deviation $\text{std}(X)$ is $\sqrt{\text{var}(X)}$ (i.e., the square root of the variance).

decrease mean net gain if that decrease also comes with an decrease in the uncertainty of returns.

Maximization of Reward-to-Variability Ratio

Stephens and Charnov [59] introduce a risk-sensitive agent model and an optimization approach that maximizes the probability of success. Consider a solitary agent that must acquire some minimal net gain \tilde{G}^T by a time $\tilde{T} \in \mathbb{R}_{\geq 0}$. Call $\tilde{\mu}$ the expectation and $\tilde{\sigma}$ the standard deviation of net gain acquired by \tilde{T} for some given behavior. The method states that the desired risk-sensitive behavior should maximize the objective

$$\frac{\tilde{\mu} - \tilde{G}^T}{\tilde{\sigma}} \tag{3.4}$$

If the net gain random variable is location-scale¹² with identical skewness for all choices of location and scale¹³, the behavior that maximizes Equation (3.4) will also minimize the probability that the net gain is less than the \tilde{G}^T threshold¹⁴. In other words, if the agent is said to be *successful* when its net gain meets or exceeds \tilde{G}^T , then the optimal behavior will maximize the probability of success¹⁵.

¹²A family of distribution functions Ω is called *location-scale* if there exists some $F \in \Omega$ such that for all $F_1 \in \Omega$, there exists a *location* $m \in \mathbb{R}$ and *scale* $s \in \mathbb{R}_{>0}$ with $F_1(x) = F((x-m)/s)$. A random variable is location-scale if its distribution comes from such a family. This idea of a two parameter family of distribution functions comes from Rothschild and Stiglitz [54], and this definition of such a class of functions is due to Bawa [6]; however, Meyer [40] gives an equivalent definition. Examples of location-scale distributions are the normal, exponential, , and double exponential distributions.

¹³Location-scale distributions with mean locations and standard deviation scales will naturally have this property.

¹⁴This is a sufficient condition; however, it is not necessary. Investment theoretic consequences of location-scale distributions are given by Bawa [6] and Meyer [40]. The multivariate case is handled by Chamberlain [13] and Owen and Rabinovitch [44].

¹⁵This result can be generalized slightly by considering the class of distributions where a monotonic transformation (i.e., continuously differentiable with non-negative derivative everywhere) of random variables is location-scale. The log-normal distribution belongs to this more general class [6].

Location-Scale Justification: By the *central limit theorem (CLT)*, if the net gain is a sum of i.i.d. random variables (e.g., individual cycle gains), the probability distribution of the net gain will approach a normal distribution¹⁶ as the number of elements in the sum increases. Therefore, it may be reasonable (e.g., consider \tilde{G}^N as $N \rightarrow \infty$) to assume that net gains are normally distributed or at least location-scale with location-scale invariant skewness. In this case, the behavior that maximizes Equation (3.4) will certainly maximize the probability of success.

Analogous Results from Economics: Stephens and Krebs [60] call this the *z-score model*; however, it is well-known in economics that this method was initially developed by Sharpe [57] for application to optimal portfolio selection. Sharpe calls Equation (3.4) a *reward-to-variability ratio*¹⁷. While economists realize that return distributions need not be normally distributed (e.g., symmetric about the mean) for the reward-to-variability ratio to minimize risk, Stephens and Krebs [60] depend on normality to justify their claims [60, p. 134]. Assuming normality of returns may be far too restrictive. In fact, it is desirable that returns are skewed so that the mass is concentrated on higher gains (i.e., not symmetric and therefore not normal). Therefore, by depending on consistent skewness rather than symmetry, the economic argument of reward-to-variability maximization is not only more general but also more convincing than the argument of Stephens and Krebs.

¹⁶A *normal* or *Gaussian* random variable X with mean μ and standard deviation σ has (Lebesgue) probability density $f_X(x) = 1/(\sigma\sqrt{2\pi}) \exp(-(x-\mu)^2/(2\sigma^2))$. Normal random variables are location-scale with location μ and scale σ and are symmetric about their mean (i.e., they have zero skewness).

¹⁷This is also known as the *Sharpe ratio*, which is named after the Nobel laureate who developed it.

Links to Risk-Sensitive Dynamic Optimization: An ex post version of the reward-to-variability ratio is described by Sharpe [58]¹⁸, which is typically used for measuring past performance. However, there may some opportunity to use this ex post ratio for dynamic optimization to derive an optimal control similar to the bang-bang¹⁹ control described by Stephens and Krebs [60] and McNamara [39] and based on the z -score model.

Graphical Interpretation of Risk Minimization

Again, consider an agent described by the model in Chapter 2 with tasks that are patchily distributed. Assume that there is a single task type and no search or processing costs (i.e., $n = 1$, $c^s = c_1 = 0$, $p_1 = 1$, and $\tau_1 \in \mathbb{R}_{\geq 0}$), but assume that the agent must acquire a net gain of \tilde{G}^T in \tilde{T} time. Stephens and Charnov [59] use this scenario to illustrate reward-to-variability maximization, as shown in Figure 3.2. The curve shows the $(\tilde{\mu}, \tilde{\sigma})$ combinations that result from each choice of processing time τ_1 , where processing times increase in a counter-clockwise direction. Because the gain threshold is so high, the optimal processing time, denoted t^* , is biased toward a higher-variance distribution of outcomes. In other words, the negative energy budget forces the agent to take more risks to maximize the probability of success. In Section 3.3.2, we demonstrate how this graphical solution can be used to gain insight into the most general case of the model from Chapter 2²⁰.

¹⁸Sharpe [58] also discusses other related ratios which may have applicability to stochastic optimization of static behavior.

¹⁹A *bang-bang* control switches an actuator from one extreme to the other. In this context, the variance at each small time period could be chosen by the agent to be either its lowest value or its highest value.

²⁰Again, while our solution will not be in terms of the statistics used in classical OFT, it can easily be extended with little more than a change of notation. Also, to maintain tractability, we assume Equations (2.35)–(2.38).




Figure 3.2: Risk sensitivity in classical OFT. Over a time period of length \tilde{T} , $\tilde{\mu}$ is the expected net gain, $\tilde{\sigma}$ is the standard deviation of the net gain, and \tilde{G}^T is an average net gain success threshold. The curve shows $(\tilde{\mu}, \tilde{\sigma})$ combinations for each average processing time, and t_a and t_b are examples of average processing times. The optimal average processing time is denoted t^* and the corresponding maximal reward-to-variability ratio is denoted by $\tilde{\rho}^*$ and shown as the slope of a tangent line.

3.2.3 Criticisms of the OFT Approach

Despite its successes in making qualitative predictions that agree with empirical data, rate maximization of the classical OFT approach is often criticized. For example, Nonacs [43] tabulates 19 cases of foragers being observed to process tasks longer than predicted by rate maximization, which shows that observations tend to deviate only in one direction. Nonacs concludes that “something fundamental is missing” [43, p. 71] from the classical OFT approach. However, as we discussed, rate maximization exerts an upward pressure on time when energy budgets are negative. Thus, the deficiency may not be in classical OFT but in gain models that do not properly include negative cases²¹. Additionally, the risk-sensitive analysis of Stephens and

²¹A generalization of this suggestion is given in Section 3.3.2.

Charnov [59] also explains this upward pressure on time for similar reasons. However, this analysis receives significantly less attention than the rate maximization approach even though analogous risk-sensitive methods are used with great success in other fields, like MPT.

Alternate approaches choose to abandon the classical OFT agent model [e.g., 4, 5, 12, 23, 26, 27, 52, 61]; these approaches study risk aversion, satisficing²², heuristic rules, and even propose fitting empirical data directly without any functional justification for the resulting behavioral rules. We value the generality and intuitive appeal of the classical OFT approach. Additionally, even though it may be debatable whether it can be used for evolutionary analysis in biology, optimization has natural applications in engineering. Therefore, we use our generalized agent model and build optimization objectives inspired by both classical OFT and its critics; these objectives should have applicability to engineering and may answer some of the questions about the shortfalls of the predictive power of classical OFT.

3.3 Generalized Optimization of Solitary Agent Behavior

Here, we discuss optimization criteria that may have utility in both engineering and biology. As discussed in [Section 3.2](#), the analysis of asymptotic behaviors may be justifiable in evolutionary analysis. However, in engineering applications, it may be more appropriate to focus on the case where an agent’s finite lifetime is determined entirely by a set number of tasks the agent must *complete*. Also, because design takes the place of natural selection, we need to explore different optimization objectives. Therefore, in this section, we use our generalized solitary agent model and insights

²²Satisficing describes a behavior that is suboptimal but achieves some minimum gain (i.e., “survival of the more fit over the less fit, not necessarily of the most fit” [12, p. 640]).

from both OFT and MPT to generate new optimization objectives that may be applicable to both biology and engineering. We discuss our approach to handling finite lifetimes in [Section 3.3.1](#). In [Section 3.3.2](#), we present functions that, like classical OFT, use rates to trade off between multiple optimization objectives. In [Section 3.3.3](#), we present generalized Pareto optimal combinations of optimization objectives that allow a design to tune the relative importance of the objectives. Note that when we discuss standard deviation in variance, we assume [Equations \(2.35\)–\(2.38\)](#) for simplicity.

Some of the objectives that we define still operate on the sliding scale of behaviors described by Schoener [\[55\]](#). The primary justification in classical OFT for using the long-term rate of gain discussed in [Section 3.2.1](#) is that it agrees with Schoener and makes predictions that agree *qualitatively* with empirical evidence [\[16, 49, 60\]](#). However, some of our alternative objectives also agree with Schoener and have qualitative predictions similar to those from classical OFT. Therefore, these criteria seem as well-suited for biological application as the long-term rate of gain. In fact, as discussed in [Section 3.3.2](#), the cases where our approaches yield different predictions than the classical OFT approach may be interesting and help explain some of the empirical inconsistencies with the long-term rate of gain approach [e.g., [4, 23, 43](#)].

3.3.1 Finite Task Processing

One disadvantage of studying behavior that is optimal with respect to the long-term rate of gain is the necessity of justifying the use of infinite-time (i.e., infinite-cycle) approximation. A true finite-time optimization approach is desirable. There are several dynamic programming approaches to time-limited agents [e.g., [18, 24, 34](#),

66] where behaviors are based on state. While this state-based approach provides can be better for the analysis animal behavior, it sometimes provides less intuition and is more difficult to implement on-line in an engineered agent that may have very limited computational abilities. Additionally, assuming that the time horizon is known or even fixed may be unrealistic. Thus, an approach that combines elements of both the infinite-time approach and the time-limited approach could be useful.

Rather than fixing the time horizon, we fix the number of tasks processed. This is useful for modeling, for example, a situation where an agent expends a limited resource on each processed task. Therefore, while the total searching and processing time is finite and unknown (i.e., random), the total number of tasks to process is fixed. So, take $N^p \in \mathbb{N}$ to be a fixed number of tasks completed by an agent. Our objective functions are based on statistics from our modeling approach. Of course, these statistics take the total number of tasks N^p as a parameter. Therefore, it makes sense that behaviors for a small number of total tasks may vary greatly from behaviors for a large number of tasks. However, in all cases behaviors will not be based on state. In particular, the behavior for the first task processed will be consistent with the behavior for the last task processed. This allows for the construction of behaviors that can be implemented on very simple engineered agents.

3.3.2 Tradeoffs as Ratios

Ratios are unique in their ability to apply pressure on one objective conditioned on the sign of another objective. However, the derivation of the expectation of ratios of random variables is not trivial and often requires high-order statistics of the random variables. Despite the drawbacks, we explore some ratios of statistics here that may

be useful when there are no other easy ways to apply the appropriate pressures for optimization. A more general method of multiobjective optimization is explored in [Section 3.3.3](#).

For example, classical OFT studies the optimization objective $E(\tilde{G}_1)/E(\tilde{T}_1)$ because it is the limiting expression of $E(\tilde{G}(t))/t$ as $t \rightarrow \infty$ and maximization of the long-term rate of gain is desirable. The analogous expression with our assumption of a finite number of tasks processed is $E(G^{N^p}/T^{N^p})$, but this expectation has poor analytical tractability. Although not ideal, maximization of $E(G_1)/E(T_1)$ provides similar pressures on behaviors and has a relatively simple analytical structure. Therefore, it is reasonable to study behaviors that are optimal with respect to this latter objective.

Rate Maximization with Gain Threshold

Consider an agent that must achieve an expected net gain of at least G^T after N^p tasks processed but must also achieve its goal in as little time as possible. This problem has aspects in common with both rate maximization and risk sensitivity. Therefore, we consider the objective function

$$\frac{E(G^{N^p}) - G^T}{E(T^{N^p})} \tag{3.5}$$

or, equivalently,

$$\frac{E(G_1) - \frac{G^T}{N^p}}{E(T_1)} \tag{3.6}$$

So, the goal of the agent is to maximize its expected *excess* net gain in as little expected time possible. This may be considered a generalization of the long-term rate of gain optimization used in classical OFT.

$$\begin{aligned}
\mu &\triangleq E(G_1) = \overline{g^p} - \overline{c^p} - \frac{c^s}{\lambda^p} \\
E(T_1) &= \overline{\tau^p} + \frac{1}{\lambda^p} \\
\gamma &\triangleq \frac{\mu - \frac{G^T}{N^p}}{\overline{\tau^p} + \frac{1}{\lambda^p}} \\
\diamond^* &\triangleq \max\{\diamond\}
\end{aligned}$$




Figure 3.3: Rate maximization. Curves show the Pareto frontiers for maximizing net gain mean and maximizing or minimizing total time. G^T denotes the net gain survival threshold. γ^* denotes the maximal rate, and the solutions that achieve this rate are shown as large dots at points of tangency with a line of slope γ^* . The dotted arrow in (b) shows how the figure changes as with c^s and λ^p . It is not necessary that $|\mu^*| < \infty$ and $(\overline{\tau^p} + 1/\lambda^p)^* < \infty$.

A graphical solution to this problem is shown in Figure 3.3. This graphical approach is similar to the time-constrained approach of Rapport [51] that is described above. The curve shows the set of all Pareto efficient solutions for maximizing gain and both maximizing or minimizing time²³. In Figure 3.3(b), the axes are shifted

²³That is, points on the curve represent the highest and lowest expected total times possible for a given expected net gain.

to highlight the impact of parameter changes on the solution²⁴. Figure 3.1 was generated in an identical way; however, the single gain curve $g_1(\tau_1)$ is replaced with a curve representing the Pareto-efficient-average-processing-gain curve. This graphical interpretation of rate maximization applies to any case rather than just the single type patch case²⁵. Additionally, search and processing costs are shown. In fact, when $G^T = 0$, it is clear that the cost of searching plays the same role as a gain success threshold. The line whose slope is the rate may be viewed as *pivoting* at the point $(-1/\lambda^p, G^T/N^p + c^s/\lambda^p)$, and so this may be called the *threshold pivot point*.

For this particular \bar{g}^p and $\bar{\tau}^p$, several observations can be made.

Cycle Lifetime and Gain Threshold: The per-cycle net gain threshold G^T/N^p depends upon the number of lifetime cycles N^p . Therefore, the length of lifetime has an impact on the gain budget and therefore the optimal strategy.

Negative Gain Budgets: Whenever the expected gain is less than the gain success threshold, the average time spent processing tasks (which the agent can control through both preference probability and processing time) will increase. Graphically, whenever $G^T/N^p + c^s/\lambda^p > (\bar{g}^p - \bar{c}^p)^*$, an increase in G^T will bring an increase in the $\bar{\tau}^p$ solution. Therefore, rate maximization in classical OFT may predict low processing times compared to observations because costs or required energy thresholds are being neglected.

²⁴The encounter rate λ^p may only be considered to be an environmental parameter if, for all $i \in \{1, 2, \dots, n\}$, p_i is not free for the agent to choose. The analogous graphical interpretation of results on classical OFT statistics does not have this constraint. However, in both cases the shape of the Pareto frontier curve will change with changes in the overall encounter rate if the individual task type encounter rates are not also scaled by the same factor.

²⁵Recall that $\bar{g}^p = g_1$ and $\bar{\tau}^p = \tau_1$ for the $n = 1$ (and $p_1 = 1$) case.

Sunk Cost Effect and Concorde Fallacy: Arkes and Ayton [3] review both the *sunk cost effect* and the *Concorde fallacy*, which are well-known in economics and biology. These describe the same behavior, except the former is typically applied to humans and the latter is typically applied to animals. Agents are said to manifest this behavior when showing a “greater tendency to continue an endeavor once an investment in” gain has been made [p. 591 3]. In other words, an agent will continue to process tasks evidently *because* of the *cost* of prior processing. Arkes and Ayton suggest that this is due to the use of an ordinarily adaptive heuristic rule; however, we suggest that rate maximization may also explain this effect. Consider the situation shown in Figure 3.3(b), but replace the strictly positive $(\overline{g^p} - \overline{c^p}) - \tau^p$ curve with a version of itself that has been mirrored around the τ^p -axis. In this case, the optimal average processing time is $(\tau^p)^*$. That is, it is better to face a negative energy budget after a long period of time rather than facing a negative energy budget after a short period of time, even though the later negative energy budget is worse (i.e., it may be better to be short in gain after much effort than being short in gain after no effort).

Therefore, rate maximization explains some effects that may otherwise be considered to be irrational.

Efficiency

Now, take an agent that must achieve an expected *gross* gain of at least G_g^T after N^p tasks processed but must achieve its goal while accumulating as little cost as possible. We consider the objective function

$$\frac{\mathbb{E}(G^{N^p}) + \mathbb{E}(C^{N^p}) - G_g^T}{\mathbb{E}(C^{N^p})} \tag{3.7}$$




Figure 3.4: Efficiency maximization. Curves show Pareto frontiers for maximizing gross gain mean and minimizing total cost. G_g^T denotes the gross gain survival threshold. ϵ^* denotes the maximal efficiency, and the solutions that achieve this efficiency are shown as large dots at points of tangency with a line of slope ϵ^* . The dotted arrow in (b) shows how the figure changes as with c^s and λ^p . It is not necessary that $|(\bar{g}^p)^*| < \infty$ and $|(\bar{c}^p + c^s/\lambda^p)^*| < \infty$.

or, equivalently,

$$\frac{E(G_1) + E(C_1) - \frac{G_g^T}{N^p}}{E(C_1)} \quad (3.8)$$

So, the goal of the agent is to maximize its expected *excess* gross gain while also accumulating as little cost as possible. This is a notion of the agent's *efficiency* (i.e., maximal benefit-to-cost ratio). While there is no direct pressure on time minimization in this objective, all costs in the model depend linearly on time. Therefore, minimization of cost indirectly has a time minimization effect as well.

A graphical solution to this problem is shown in [Figure 3.4](#). Here, the curve shows

the set of all Pareto efficient solutions for maximizing gross gain and both maximizing or minimizing cost²⁶. In Figure 3.4(b), the axes are shifted to highlight the impact of parameter changes on the solution²⁷. The line whose slope is the efficiency may be viewed as *pivoting* at the point $(-c^s/\lambda^p, G_g^T/N^p)$, and so this may be called the *threshold pivot point*.

For this particular \bar{g}^p and \bar{c}^p , several observations can be made.

Cycle Lifetime and Gain Threshold: As with rate maximization, the per-cycle gross gain threshold G_g^T/N^p depends upon the number of lifetime cycles N^p . Therefore, the length of lifetime has an impact on the gain budget and the optimal strategy.

Negative Gain Budgets: Assuming the gain threshold is high, the optimal behavior moves toward has a higher cost and a lesser gross gain as the gross gain threshold is increased. In other words, the number of points *lost* per unit cost (i.e., the loss efficiency) is actually lower at a higher cost.

Negative Search Costs: Assume that $c^s < 0$ and $G_g^T \leq 0$. In this case, the optimal solution is the one such that $-c^s/\lambda = \bar{c}^p$. In other words, when facing a positive energy budget and a search *gain*, the most efficient solution is the one where search gains are equal to processing costs. A similar result holds for negative average search costs. Essentially, whenever the gain threshold pivot point is

²⁶That is, points on the curve represent the highest and lowest expected cost possible for a given expected gross gain.

²⁷As with the graphical interpretation of rate maximization, the effects of parameter changes are simpler to explore using classical OFT statistics because the overall encounter rate is not influenced by the preference probabilities, which are often considered to be decision variables that describe agent behavior.

below (or encircled by) the Pareto frontiers, the optimal solution will be the one that makes $E(C_1) = 0$.

Sunk Cost Effect and Concorde Fallacy: Efficiency maximization provides an explanation for the sunk cost effect and the Concorde fallacy similar to the explanation for rate maximization. If the \bar{g}^p - \bar{c}^p curve in Figure 3.4(b) is mirrored about the \bar{c}^p -axis, the resulting optimal cost will be $(\bar{c}^p)^*$. That is, processing cost is maximized in this case. Roughly, it is “better” to have a negative gain budget at high cost than low cost.

As with rate maximization, an analysis of this optimization objective suggests explanations for behaviors that might normally be considered irrational.

Risk-Sensitivity: Reward-to-Variability Ratios

As discussed, Sharpe [57] introduces an index called the *reward-to-variability ratio*²⁸, a measure of future performance that balances expected return with standard deviation of return. When returns are location-scale distributed with identical skewness for each location-scale choice (e.g., normally distributed), maximization of this index is identical to minimization of risk of return shortfall.

Let G^T represent a (deterministic) net gain threshold of success for the N^p cycles in an agent’s lifetime. In our context²⁹, the reward-to-variability ratio is expressed by

$$\frac{E(G^{N^p}) - G^T}{\text{std}(G^{N^p})} \quad \text{or, equivalently,} \quad \sqrt{N^p} \frac{E(G_1) - \frac{G^T}{N^p}}{\text{std}(G_1)}$$

That is, this is a ratio of *excess* returns to the standard deviation of those returns.

²⁸Again, this is often called the *Sharpe ratio* by those other than Sharpe.

²⁹For simplicity, we assume Equations (2.35)–(2.38).

A graphical solution to this problem is shown in Figure 3.5. The curve represents the Pareto frontiers for maximizing expected net gain and minimizing or maximizing net gain standard deviation. Notice that the shape of the curve shown in Figure 3.5(a) differs slightly from the shape of the curve in Figure 3.5(b) due to different scaling of the axes. Even when λ^p may be considered a parameter of the environment (i.e., not a function of decision variables), the shape of the Pareto frontier curve is still dependent upon it. Therefore, the graphical interpretation of maximization in Figure 3.5(c) requires horizontally squeezing or stretching the curve with changes in λ^p .

For this particular \bar{g}^p and \bar{c}^p , several observations can be made.

Cycle Lifetime and Gain Threshold: Similarly as with rate and efficiency maximization, the per-cycle gross gain threshold G^T/N^p depends upon the number of lifetime cycles N^p . Therefore, the length of lifetime has an impact on the gain budget and therefore the optimal strategy.

Negative Gain Budgets: Assuming the gain threshold is high, the optimal behavior moves toward has a higher standard deviation and a lesser net gain as the net gain threshold is increased. In other words, the number of points *lost* per unit standard deviation is actually lower at a higher standard deviation. Put another way, additional uncertainty in rewards is less costly because it comes with a higher probability that returns will be above the gain threshold.

As with rate and efficiency maximization, an analysis of this optimization objective suggests explanations for behaviors that might normally be considered irrational.

$$\begin{aligned} \mu &\triangleq \mathbb{E}(G_1) = \overline{g^p} - \overline{c^p} - \frac{c^s}{\lambda^p} \\ \sigma &\triangleq \text{std}(G_1) = \sqrt{\text{var}(g^p - c^p) + \left(\frac{c^s}{\lambda^p}\right)^2} \\ \rho &\triangleq \frac{\mu - G^T}{\sigma} \\ \diamond^* &\triangleq \max\{\diamond\} \end{aligned}$$

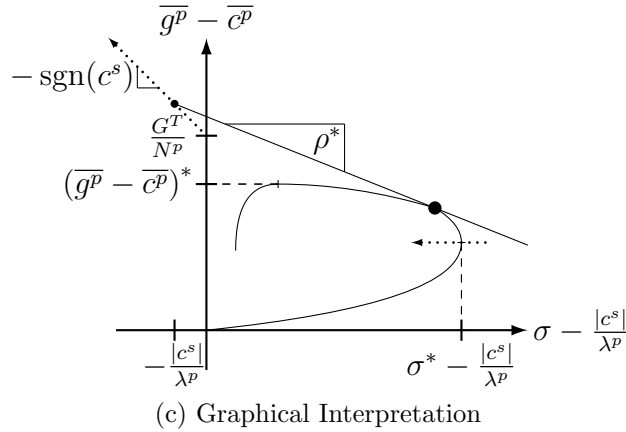



Figure 3.5: Reward-to-variability maximization. Curves show Pareto frontiers for maximizing mean net gain and minimizing net gain scale. G^T denotes the net gain survival threshold. ρ^* denotes the maximal per-cycle reward-to-variability ratio, and the solutions that achieve this ratio are shown as large dots at points of tangency with a line of slope ρ^* . The dotted arrows in (c) show how the figure changes as $|c^s|/\lambda^p$ increases. It is not necessary that $|\mu^*| < \infty$ and $\sigma^* < \infty$.

Reward-to-Variance Ratios

Optimization of portfolio performance based entirely on expectation and variance was originally developed by Markowitz [36] and Tobin [62]. Sharpe [57] makes the natural substitution of standard deviation for variance in the definition of the reward-to-variability ratio. While this substitution allows maximization of the performance index to have a risk minimization interpretation, it makes analytical derivation of the optimal behavior difficult. To compensate, we introduce a *reward-to-variance* ratio that is motivated by the reward-to-variability ratio but provides greater analytical tractability.

Let G^T represent a (deterministic) net gain threshold of success for the N^p cycles in an agent's lifetime. In our context³⁰, the reward-to-variance ratio is expressed by


$$\frac{E(G^{N^p}) - G^T}{\text{var}(G^{N^p})} \quad \text{or, equivalently,} \quad N^p \frac{E(G_1) - \frac{G^T}{N^p}}{\text{var}(G_1)}$$

That is, this is a ratio of *excess* returns to the *variance* of those returns.


A graphical solution to this problem is shown in [Figure 3.6](#). The curve represents the Pareto frontiers for maximizing expected net gain and minimizing or maximizing net gain variance. Notice that the shape of the the curve shown in [Figure 3.6\(a\)](#) is identical to the shape of the curve in [Figure 3.6\(b\)](#), which relates to the analytical tractability gained by using the reward-to-variance ratio over the reward-to-variability ratio. Observations made about the reward-to-variability ratio also hold for the reward-to-variance ratio, and so we omit a discussion of them here.

³⁰For simplicity, we assume [Equations \(2.35\)–\(2.38\)](#).

$$\begin{aligned} \mu &\triangleq E(G_1) = \bar{g}^p - \bar{c}^p - \frac{c^s}{\lambda^p} \\ \sigma^2 &\triangleq \text{var}(G_1) = \text{var}(g^p - c^p) + \left(\frac{c^s}{\lambda^p}\right)^2 \\ \eta &\triangleq \frac{\mu - \frac{G^T}{N^p}}{\sigma^2} \\ \diamond^* &\triangleq \max\{\diamond\} \end{aligned}$$



(a) Lifetime Perspective (b) Per-Cycle Perspective



(c) Graphical Interpretation

Figure 3.6: Reward-to-variance maximization. Curves show Pareto frontiers for maximizing net gain mean and minimizing net gain variance. G^T denotes the net gain survival threshold. η^* denotes the maximal reward-to-variance ratio, and the solutions that achieve this ratio are shown as large dots at points of tangency with a line of slope η^* . The dotted parabolic arrow in (c) shows how the figure changes as $|c^s|/\lambda^p$ increases. It is not necessary that $|\mu^*| < \infty$ and $(\sigma^2)^* < \infty$.

3.3.3 Generalized Pareto Tradeoffs

As discussed, maximization of a rate has the ability to choose solutions that are Pareto efficient with respect to maximization of one objective and maximization or minimization of another objective, conditioned on the sign of the first objective. The particular Pareto efficient solution chosen depends on the shapes of the optimization functions. In other words, the relative importance of maximizing one objective and minimizing another objective is allowed to vary. Therefore, maximization of a rate is equivalent to picking a particular relative importance.

In some cases, it may be useful to let the relative importance of maximization of one objective and minimization of a different objective be a parameter of the environment. For example, for objectives A and B and weight $w \in \mathbb{R}$, the maximization of $A - wB$ produces a Pareto efficient solution to maximization of A and minimization of B , given that w is the relative importance of minimizing B over maximizing A . If $w < 0$, then $|w|$ represents the relative importance of *maximizing* B over maximizing A . Maximization of $A - wB$ is sufficient but not necessary for a solution to be Pareto efficient with respect to objectives A and B ; therefore, unless A and B are convex functions, there may exist additional Pareto efficient solutions that are not accessible by this method. Either way, collecting solutions to this maximization problem for all $w \in \mathbb{R}$ is one way of generating a Pareto efficient set.

Therefore, here we recast the tradeoffs from [Section 3.3.2](#) as optimization problems of relative importance. Keeping in mind that maximization of $A - wB$ is a sufficient but not necessary condition for Pareto efficient solutions, this process can be viewed as picking a particular point on the curves shown in [Figures 3.3–3.6](#).

Gain and Time

Take some weight $w \in \mathbb{R}$ that represents the relative importance of decreasing total time over increasing net gain (i.e., if $w < 0$, then $|w|$ is the relative importance of *increasing* total time over increasing total net gain). Solutions that maximize

$$E(G_1) - w E(T_1) \tag{3.9}$$

will be Pareto efficient with respect to these two objectives. If all of the solutions for all $w \in \mathbb{R}$ are collected, the result will be a curve like the one shown in [Figure 3.3](#). This is analogous to optimizing $E(\tilde{G}_1) - w E(\tilde{T}_1)$ in a classical OFT context. Schoener [55] considered net gain maximization and time minimization as two extremes on a continuum of behaviors that would maximize reproductive success in a certain environment. If $w \in \mathbb{R}_{\geq 0}$, then this problem captures this idea, where w picks a particular Pareto efficient solution for a certain environment. As discussed by Houston and McNamara [24], when w is set to the maximum long-term rate of gain (i.e., the maximum value of $E(G_1)/E(T_1)$), the behavior that maximizes [Equation \(3.9\)](#) will also be the behavior that maximizes the long-term rate of gain.

Efficiency

Take some weight $w \in \mathbb{R}$ that represents the relative importance of decreasing total cost over increasing *gross* gain (i.e., if $w < 0$, then $|w|$ is the relative importance of *increasing* total cost over increasing total net gain). Solutions that maximize

$$E(G_1 + C_1) - w E(C_1) \quad \text{or, equivalently,} \quad E(G_1) + (1 - w) E(C_1) \tag{3.10}$$

will be Pareto efficient with respect to these two objectives. If all of the solutions for all $w \in \mathbb{R}$ are collected, the result will be a curve like the one shown in [Figure 3.4](#).

Risk-Sensitivity: Mean and Standard Deviation

Take some weight $w \in \mathbb{R}$ that represents the relative importance of decreasing net gain standard deviation³¹ over increasing expected net gain (i.e., if $w < 0$, then $|w|$ is the relative importance of increasing net gain standard deviation over increasing expected net gain). Solutions that maximize

$$E(G^{N^p}) - w \text{std}(G^{N^p}) \quad \text{or, equivalently,} \quad E(G_1) - \frac{w}{\sqrt{N^p}} \text{std}(G_1) \quad (3.11)$$

will be Pareto efficient with respect to these two objectives. If all of the solutions for all $w \in \mathbb{R}$ are collected, the result will be a curve like the one shown in [Figure 3.5](#).

Risk-Sensitivity: Mean, Variance, and Expected Utility

The analytical tractability of [Equation \(3.11\)](#) can be low. However, substituting variance for standard deviation achieves a similar goal while allowing for a simpler analysis. That is, take some weight $w \in \mathbb{R}$ that represents the relative importance of decreasing net gain variance over increasing expected net gain (i.e., if $w < 0$, then $|w|$ is the relative importance of increasing net gain variance over increasing expected net gain). Solutions that maximize

$$E(G^{N^p}) - w \text{var}(G^{N^p}) \quad \text{or, equivalently,} \quad E(G_1) - w \text{var}(G_1) \quad (3.12)$$

will be Pareto efficient with respect to these two objectives. If all of the solutions for all $w \in \mathbb{R}$ are collected, the result will be a curve like the one shown in [Figure 3.6](#).

This approach was suggested by Real [52] as a way of applying the mean-variance methods of Markowitz [36] and Tobin [62] to biology. Real calls this *variance discounting*. The parameter w reflects a shape parameter of the agent's utility function³².

³¹Again, for simplicity, we assume [Equations \(2.35\)–\(2.38\)](#).

³²A *utility function* quantifies the preferences of an agent. For example, a utility function shaped one way may indicate that an agent is risk prone whereas a utility function shaped another way

Real uses an approximation of a general utility function to argue that maximizing Equation (3.12) leads to maximal expected utility. Stephens and Krebs [60] make this argument exact by using an approach of Caraco [11] and some assumptions about the shape of the utility function and the distribution of returns. In either case, there must be some justification for the choice of w ; information needs to be known about the shape of the agent's utility function. In an engineering context, this means that w is a parameter of a utility function that captures the design preferences. The optimization of Equation (3.12) should on average lead to the most preferable outcome with respect to this utility function³³.

3.3.4 Constraints

In Sections 3.3.2 and 3.3.3, multiple optimization objectives are traded off in a way to combine elements of both. In order to introduce threshold effects, the optimization problems are ultimately formulated in terms of *excess*. These thresholds could be made more strict by acting on the means of the distributions rather than on the outcomes. Thus, for some threshold \hat{X} , instead of maximizing $X - \hat{X}$, we can maximize X subject to the constraint that $E(X) \geq \hat{X}$. These type of constraints can be viewed as drawing horizontal and vertical boundaries on the Pareto frontiers in Figures 3.3–3.6. The optimal solution remains Pareto efficient with respect to may indicate that an agent is risk avoiding. Economists conventionally assumed that agents make decisions that maximize future expected utility. The foundations of modern utility theory are due to von Neumann and Morgenstern [65]; these results are summarized by Marschak [38]. Utility theory will be discussed further in Section 3.4.2.

³³The ratio-based risk-sensitivity methods in Section 3.3.2 may be viewed as choosing an agent's utility function and therefore its preferences based on its environment. Here, the preferences are a design variable.

all objectives in question; however, it will be optimal with respect to one particular objective when defined over the constrained domain. Unfortunately, analytically defining this constrained domain may be challenging.

Gain and Time

Let T be some threshold of time. Consider the objective

$$\text{maximize } E(G^{N^p}) \quad \text{subject to} \quad E(T^{N^p}) \leq T$$

This is identical to

$$\text{maximize } E(G_1) \quad \text{subject to} \quad E(T_1) \leq \frac{T}{N^p}$$

This is equivalent to drawing a vertical line at $E(T_1) = T/N^p$ on [Figure 3.3\(a\)](#) and choosing the point on the Pareto efficient curve to the left of that line that has the highest $E(G_1)$. Notice that as N^p increases, the feasible set of behaviors decreases, which eventually drives down the maximum possible net gain.

Similarly, let G^T be some threshold of net gain. Consider the objective

$$\text{minimize } E(T^{N^p}) \quad \text{subject to} \quad E(G^{N^p}) \geq G^T$$

This is identical to

$$\text{minimize } E(T_1) \quad \text{subject to} \quad E(G_1) \geq \frac{G^T}{N^p}$$

This is equivalent to drawing a horizontal line at $E(G_1) = G^T/N^p$ on [Figure 3.3\(a\)](#) and choosing the point on the Pareto efficient curve above that line that has the lowest $E(T_1)$. Notice that as N^p increases, the feasible set of behaviors increases, which eventually deactivates the constraint³⁴.

³⁴Borrowing language from Lagrange multiplier methods [10], a constraint is not *active* when it is satisfied with strict inequality. Constraints form boundaries around solutions; when a solution falls on a boundary, the constraint is said to be active, which implies that the optimal solution would most likely fall outside of the boundary if it were removed.

Efficiency

Let C^T be some cost threshold. Consider the objective

$$\text{maximize } E(G^{N^p}) + E(C^{N^p}) \quad \text{subject to} \quad E(C^{N^p}) \leq C^T$$

This is identical to

$$\text{maximize } \bar{g}^p \quad \text{subject to} \quad E(C_1) \leq \frac{C^T}{N^p}$$

This is equivalent to drawing a vertical line at $E(C_1) = C^T/N^p$ on [Figure 3.4\(a\)](#) and choosing the point on the Pareto efficient curve to the left of that line that has the highest \bar{g}^p . Notice that as N^p increases, the feasible set of behaviors decreases, which eventually drives down the maximum possible gross gain.

Similarly, let G_g^T be some threshold of *gross* gain. Consider the objective

$$\text{minimize } E(C^{N^p}) \quad \text{subject to} \quad E(G^{N^p}) + E(C^{N^p}) \geq G_g^T$$

This is identical to

$$\text{minimize } E(C_1) \quad \text{subject to} \quad \bar{g}^p \geq \frac{G_g^T}{N^p}$$

This is equivalent to drawing a horizontal line at $\bar{g}^p = G_g^T/N^p$ on [Figure 3.4\(a\)](#) and choosing the point on the Pareto efficient curve above that line that has the lowest $E(C_1)$. Notice that as N^p increases, the feasible set of behaviors increases, which eventually deactivates the constraint.

Risk-Sensitivity: Mean, Standard Deviation, and Variance

Let $\hat{\sigma}^2$ be some threshold of net gain variance. Consider the objective

$$\text{maximize } E(G^{N^p}) \quad \text{subject to} \quad \text{var}(G^{N^p}) \leq \hat{\sigma}^2$$

This is identical to

$$\text{maximize } E(G_1) \quad \text{subject to} \quad \text{var}(G_1) \leq \frac{\hat{\sigma}^2}{N^p}$$

This is equivalent to drawing a vertical line at $\text{var}(T_1) = \hat{\sigma}^2/N^p$ on [Figure 3.6\(b\)](#) and choosing the point on the Pareto efficient curve to the left of that line that has the highest $E(G_1)$. Notice that as N^p increases, the feasible set of behaviors decreases, which eventually drives down the maximum possible net gain.

Similarly, let G^T be some threshold of net gain. Consider the objective

$$\text{minimize } \text{var}(G^{N^p}) \quad \text{subject to} \quad E(G^{N^p}) \geq G^T$$

This is identical to

$$\text{minimize } \text{var}(G_1) \quad \text{subject to} \quad E(G_1) \geq \frac{G^T}{N^p}$$

This is equivalent to drawing a horizontal line at $E(G_1) = G^T/N^p$ on [Figure 3.6\(b\)](#) and choosing the point on the Pareto efficient curve above that line that has the lowest $\text{var}(G_1)$. Notice that as N^p increases, the feasible set of behaviors increases, which eventually deactivates the constraint.

Standard Deviation: The optimization objectives here could be rewritten using standard deviation. That is, variance could be replaced with standard deviation and any constraints on variance could be replaced with its square root. The resulting optimization objective leads to the same solutions as the original objective. Therefore, we omit a special standard deviation case. Often, the use of variance will improve analytical tractability of these problems.

3.4 Future Directions Inspired by PMPT

As with classical OFT, results here are influenced primarily by MPT, which describes a set of methods for portfolio investment and capital budgeting that follows from the work of Markowitz [36, 37] and Tobin [62] and applications of that work by Lintner [33], Mossin [41], and Sharpe [56]. MPT is mostly concerned with mean-variance analysis (MVA). Proponents of MPT openly acknowledge that MVA is often too naive [37, 56]; however, its use is justified because it is well-understood and has low computational demands. However, modern technological and theoretical advances may render these justifications invalid [53]. Thus, PMPT seeks to improve upon MPT by using more advanced analytical approaches. We discuss two PMPT topics here that may have potential applications in biology and engineering in the analysis and design of agents.

3.4.1 Lower Partial Moments

We have used the well-known constructs of standard deviation and variance to quantify the variability of a return. However, for a given gain threshold, minimization of the variation *below* that threshold is much more important than the variation above the threshold. That is, rather than seeking certainty in future gains, it may be more useful to minimize the uncertainty in negative gain budgets. Therefore, Markowitz [37] introduces the *semivariance* which Bawa [6] calls the *lower-partial variance (LPV)*. The LPV is the expected value of the squared negative deviations of possible outcomes from some point of reference. That is, for a distribution F and a reference point t , the LPV is defined by

$$\text{LPV}(t; F) \triangleq \int_x^t (t - x)^2 dF(x)$$

where the integral is the Lebesgue integral. Bawa and Lindenberg [9] generalize this idea with the notion of an n^{th} order *lower-partial moment* (*LPM*) defined by

$$\text{LPM}_n(t; F) \triangleq \int_x^t (t - x)^n dF(x)$$

Clearly, $\text{LPV}(t; F) = \text{LPM}_2(t; F)$ for all $t \in \mathbb{R}$. These *lower-partial* moments³⁵ are generalized asymmetric notions of the familiar central moments (e.g., variance and skewness). Rather than integrating (i.e., summing) all variations around the mean of a distribution, the LPM integrates all variations that fall beneath some arbitrary reference which might be viewed as a target benchmark (e.g., a net gain threshold). This is a type of *downside risk* measure [22, 53].

The n^{th} order LPM can be substituted for variance in MVA to yield a new method of analysis that trades off greater expected returns and shortfall uncertainty without putting any pressure on variances above return benchmarks. This is called *mean-lower-partial-moment* (*MLPM*) analysis [7, 32]. When $n = 2$, this is called *mean-lower-partial-variance* (*MLPV*) analysis [9] or *mean-semivariance analysis* (*MSA*) [35]. When MLPV analysis is applied to normal distributions, the results are identical to MVA³⁶. In fact, MLPM analysis will always have results that “do at least as well” as MVA [9]. Thus, Mao [35] claims that “the time has come to shift” from MVA to MSA in capital budgeting. We believe that optimization based on these new downside risk frameworks may also lead to advances in behavioral analysis and solitary agent design.

³⁵Bawa [7] shows that for all $n \in \mathbb{N}$ and $t \in \mathbb{R}$, $\text{LPM}_n(t; F)$ is a convex function of F .

³⁶This is due to the symmetry of normal distributions.

3.4.2 Stochastic Dominance

Investment returns (e.g., net gain) are naturally random variables. Thus, portfolios or budgets (e.g., processing probability and time choices) correspond to random variables with different probability distributions. Portfolio selection is thus a decision problem of choosing the most preferable probability distribution. von Neumann and Morgenstern [65] characterize preferences in terms of utility functions, which represent the effective value of certain returns to an agent. Different-shaped utility functions correspond to different preferences, and rational agents³⁷ will maximize their expected utility. That is, if F and G are two distributions of returns and u is an agent's utility function, the distribution F is preferred over the distribution G if

$$\int_{-\infty}^{\infty} u(x) dF(x) - \int_{-\infty}^{\infty} u(x) dG(x) > 0 \quad (3.13)$$

This rule can be used to pick the best distribution(s) from a set of distributions that correspond to portfolio choices. While each of these probability distributions may be known or at least knowable, it may not be possible to completely describe an agent's utility function. Therefore, it is desirable to define rules like Equation (3.13) that are a function of *only* the probability distributions. Distributions that are preferable with respect to these rules will return maximal utility for *all* utility functions that share certain very general characteristics. These are known as *stochastic dominance (SD)* rules. For example, as shown by Bawa [6], for any two distributions F and G , F is preferred (i.e., results in a greater expected utility) to G for any increasing and continuously differentiable utility function $u(x)$ assumed to have finite values for

³⁷von Neumann and Morgenstern [65] define what *rational agent* means with a number of axioms about preference.

finite values of x if and only if there exists some $x_0 \in \mathbb{R}$ such that

$$F(x) \leq G(x) \text{ for all } x \in \mathbb{R} \quad \text{and} \quad F(x_0) < G(x_0) \quad (3.14)$$

Equation (3.14) is called *first-order stochastic dominance (FSD)*. FSD implies that for any return benchmark, the probability of failing to meet that benchmark will be less with the F distribution³⁸. This class of utility functions for FSD is very broad. There are higher-order SD rules that apply to smaller sets of utility functions (e.g., risk averse utility functions). In particular, Bawa [6] shows that *third-order stochastic dominance (TSD)* has an important relationship to MLPM analysis. An extensive categorized bibliography of articles on SD and its application is given by Bawa [8], who also briefly discusses the foundations of SD.

Biologists have recognized the validity of utility functions [e.g., 11, 52, 60]. However, the methods of SD have not been applied to compare alternative behaviors in terms of maximal expected utility in a convincing way. That is, MVA is used to approximate an SD analysis; however, there seems to be little recognition that this approximation is only valid for a limited set of utility functions. Bawa [6] shows that for a large set of utility functions, MVA meets neither necessary nor sufficient conditions for optimality and MLPV analysis should be used instead. As design preferences are often rational, engineering can also benefit from an SD approach. If nothing else, SD provides a deeper understanding of the consequences of behavioral choice by avoiding heuristic arguments and approximations. Additionally, using SD promises to allow both engineers and behavioral ecologists to benefit not only from classical works in economics but also from the modern-day research in the field.

³⁸Of course, FSD provides no guarantee that a particular outcome from the G distribution will not succeed when the outcome from the F distribution will fail. Dominance for all outcomes is the strongest form of SD (and is rarely used).

CHAPTER 4

FINITE-LIFETIME OPTIMIZATION RESULTS

In [Chapter 3](#), we discussed selecting agent behavior based on the optimization of functions that encapsulated a number of different optimization objectives. Some of these functions traded off objectives by maximizing the ratio of one to the other. Other functions traded off objectives by maximizing a linear combination of objectives. In either case, for many of the statistics defined in [Chapter 2](#), the resulting functions have a special structure in common. In this chapter, we optimize a general value function that also has this structure. We then apply these general results to value functions of interest in our model. The general results are given in [Section 4.1](#) and the specific results are given in [Section 4.2](#).

4.1 Optimization of a Rational Objective Function

Before we discuss optimization of some of the functions introduced in [Chapter 3](#), we focus on a special general case. This general case may be applied to the optimization functions we have introduced or be used to derive optimal behavior for other novel valuation functions that have this structure.

4.1.1 The Generalized Problem

Take $n \in \mathbb{N}$ task types. Most statistics defined in [Chapter 2](#) and optimization functions described in [Chapter 3](#) have a special structure in common. Here, we present a generalized optimization problem that provides solutions to a broad range of problems in this model.

The Decision Variables and Constraints

The decision variables are preference probabilities and processing times. These variables are constrained, so their bounds must be defined as parameters. For each $i \in \{1, 2, \dots, n\}$, define upper and lower preference constraint parameters $p_i^-, p_i^+ \in [0, 1]$ and upper and lower time constraint parameters $\tau_i^- \in \mathbb{R}_{\geq 0}$, and $\tau_i^+ \in \overline{\mathbb{R}}_{\geq 0}$. Collect these constraint parameters into vectors $p^-, p^+ \in [0, 1]^n$, $\tau^- \in \mathbb{R}_{\geq 0}^n$, and $\tau^+ \in \overline{\mathbb{R}}_{\geq 0}^n$ defined by

$$\begin{aligned} p^- &\triangleq [p_1^-, p_2^-, \dots, p_n^-]^\top & \tau^- &\triangleq [\tau_1^-, \tau_2^-, \dots, \tau_n^-]^\top \\ p^+ &\triangleq [p_1^+, p_2^+, \dots, p_n^+]^\top & \tau^+ &\triangleq [\tau_1^+, \tau_2^+, \dots, \tau_n^+]^\top \end{aligned}$$

Therefore, for an arbitrary preference probability vector p and processing time vector τ defined so that

$$p \triangleq [p_1, p_2, \dots, p_n]^\top \quad \tau \triangleq [\tau_1, \tau_2, \dots, \tau_n]^\top$$

it must be that

$$p_i^- \leq p_i \leq p_i^+ \quad \text{and} \quad \tau_i^- \leq \tau_i \leq \tau_i^+$$

for all $i \in \{1, 2, \dots, n\}$.

Generalized Advantage, Disadvantage, and Objective

For each $i \in \{1, 2, \dots, n\}$, define the generalized task advantage function $a_i : \mathbb{R}_{\geq 0} \cap [\tau_i^-, \tau_i^+] \mapsto \mathbb{R}$ and the generalized task disadvantage function $d_i : \mathbb{R}_{\geq 0} \cap [\tau_i^-, \tau_i^+] \mapsto \mathbb{R}$ to be continuously differentiable functions. Also define the environment advantage $a \in \mathbb{R}$ and disadvantage $d \in \mathbb{R}$. Therefore, the total advantage A and total disadvantage D are defined by

$$A(p, \tau) \triangleq a + \sum_{i=1}^n p_i a_i(\tau_i) \qquad D(p, \tau) \triangleq d + \sum_{i=1}^n p_i d_i(\tau_i)$$

where $p \in [0, 1]^n$ and $\tau \in \mathbb{R}_{\geq 0}^n$ are arbitrary preference probability and processing time vectors. Therefore, the generalized objective J , the advantage-to-disadvantage ratio, is defined by $J(p, \tau) \triangleq A(p, \tau)/D(p, \tau)$.

Notation

Take $i, j \in \{1, 2, \dots, n\}$. For the advantage a_i and disadvantage d_i , use the notation

$$\begin{aligned} a'(\tau_i) &\triangleq \frac{d}{d\tau_i} a(\tau_i) & a''(\tau_i) &\triangleq \frac{d^2}{d\tau_i^2} a(\tau_i) \\ d'(\tau_i) &\triangleq \frac{d}{d\tau_i} d(\tau_i) & d''(\tau_i) &\triangleq \frac{d^2}{d\tau_i^2} d(\tau_i) \end{aligned}$$

to represent the first and second derivatives of each function evaluated at the point τ_i .

4.1.2 The Optimization Procedure

The goal is to choose preference probabilities and processing times to (locally) maximize J . This can be formulated as the constrained minimization problem

minimize $-J$

subject to $-\tau_i \leq \tau_i^-, \tau_i \leq \tau_i^+, -p_i \leq p_i^-, p_i \leq p_i^+$ for all $i \in \{1, \dots, n\}$

with $4n$ inequality constraints. This problem can be solved using Lagrange multiplier theory [10]. Define the Lagrangian L by

$$L \triangleq -J - \mu_-^\top(p - p^-) + \mu_+^\top(p - p^+) - \nu_-^\top(\tau - \tau^-) + \nu_+^\top(\tau - \tau^+)$$

where $\mu_-, \mu_+, \nu_-, \nu_+ \in \mathbb{R}_{\geq 0}^n$ are vectors of Lagrange multipliers, denoted

$$\begin{aligned} \mu_- &\triangleq [\mu_{1-} \ \mu_{2-} \ \dots \ \mu_{n-}]^\top & \mu_+ &\triangleq [\mu_{1+} \ \mu_{2+} \ \dots \ \mu_{n+}]^\top \\ \nu_- &\triangleq [\nu_{1-} \ \nu_{2-} \ \dots \ \nu_{n-}]^\top & \nu_+ &\triangleq [\nu_{1+} \ \nu_{2+} \ \dots \ \nu_{n+}]^\top \end{aligned}$$

For ease of notation, we use the symbol m^* to represent a collection of one of each of these four Lagrange multiplier vectors. That is, $m^* \in (\mathbb{R}_{\geq 0}^n)^4$ with $m^* \triangleq (\mu_-^*, \mu_+^*, \nu_-^*, \nu_+^*)$. Next, denote the feasible set \mathcal{F} of decision variables by

$$\mathcal{F} \triangleq \{(p, \tau) \in [0, 1]^n \times \mathbb{R}_{\geq 0}^n : p_i^- \leq p_i \leq p_i^+, \tau_i^- \leq \tau_i \leq \tau_i^+, i \in \{1, 2, \dots, n\}\}$$

Also, for each $(p^*, \tau^*) \in \mathcal{F}$, define the sets of active inequality constraints¹ $\mathcal{A}_p^-(p^*)$, $\mathcal{A}_p^+(p^*)$, $\mathcal{A}_\tau^-(\tau^*)$, $\mathcal{A}_\tau^+(\tau^*)$ by

$$\begin{aligned} \mathcal{A}_p^-(p^*) &\triangleq \{i \in \{1, 2, \dots, n\} : p_i^* = p_i^-\} & \mathcal{A}_\tau^-(\tau^*) &\triangleq \{i \in \{1, 2, \dots, n\} : \tau_i^* = \tau_i^-\} \\ \mathcal{A}_p^+(p^*) &\triangleq \{i \in \{1, 2, \dots, n\} : p_i^* = p_i^+\} & \mathcal{A}_\tau^+(\tau^*) &\triangleq \{i \in \{1, 2, \dots, n\} : \tau_i^* = \tau_i^+\} \end{aligned}$$

For any $i \in \{1, 2, \dots, n\}$, if $p_i^+ = p_i^-$ ($\tau_i^+ = \tau_i^-$), then the inequality Lagrange multipliers μ_i^+ and μ_i^- (ν_i^+ and ν_i^-) combine to form an equality Lagrange multiplier $\mu_i^+ - \mu_i^- \in \mathbb{R}$ ($\tau_i^+ - \tau_i^- \in \mathbb{R}$). Therefore, it is clear that all points $(p, \tau) \in \mathcal{F}$ are

¹An *active* inequality constraint is a constraint that holds only by equality. For example, the constraint $x \geq 1$ is active for $x = 1$ and *inactive* for $x > 1$.

regular². Finally, for any point (p^*, τ^*) , define the *feasible variations* $\mathcal{V}(p^*, \tau^*)$ by

$$\mathcal{V}(p^*, \tau^*) \triangleq \left\{ \begin{array}{l} \begin{bmatrix} \delta_1^p \\ \delta_2^p \\ \vdots \\ \delta_n^p \\ \delta_1^\tau \\ \delta_2^\tau \\ \vdots \\ \delta_n^\tau \end{bmatrix} \in \mathbb{R}^{2n} : \begin{array}{l} \delta_i^p = 0, \quad i \in \mathcal{A}_p^-(p^*) \cup \mathcal{A}_p^+(p^*), \\ \delta_j^\tau = 0, \quad j \in \mathcal{A}_\tau^-(\tau^*) \cup \mathcal{A}_\tau^+(\tau^*) \end{array} \end{array} \right\}$$

We also define the gradient operator ∇ and the Hessian operator ∇^2 by

$$\nabla \triangleq \left[\frac{\partial}{\partial p_1}, \frac{\partial}{\partial p_2}, \dots, \frac{\partial}{\partial p_n}, \frac{\partial}{\partial \tau_1}, \frac{\partial}{\partial \tau_2}, \dots, \frac{\partial}{\partial \tau_n} \right]^\top \quad \nabla^2 \triangleq \nabla \nabla^\top$$

so that we have the gradient ∇L and the Hessian $\nabla^2 L$. When these are to be evaluated at a point $(p^*, \tau^*) \in \mathcal{F}$ with multipliers $m^* \in (\mathbb{R}_{\geq 0}^n)^4$, we use the notation $\nabla L(p^*, \tau^*, m^*)$ and $\nabla^2 L(p^*, \tau^*, m^*)$, respectively. Because the Lagrangian is a continuous function, its Hessian matrix will be symmetric.

First-Order Necessary Conditions

Assume that the point $(p^*, \tau^*) \in \mathcal{F}$ is a local maximum of the objective function.

For convenience, use the notation

$$J^* \triangleq J(p^*, \tau^*) \quad A^* \triangleq A(p^*, \tau^*) \quad D^* \triangleq D(p^*, \tau^*) \quad (4.1)$$

In order for J^* to be well-defined, it must be assumed that D^* is nonzero³. It is necessary that there exist Lagrange multiplier vectors $m^* \in (\mathbb{R}_{\geq 0}^n)^4$ such that

$$\nabla L(p^*, \tau^*, m^*) = 0 \quad (4.2)$$

²In this context, a *regular* point is a point where all active constraint gradients are linearly independent.

³While J^* is not defined for $D^* = 0$, any case where $A^* > 0$ and $D^* = 0$ is certainly desirable.

and for all $i \in \{1, 2, \dots, n\}$,

$$i \notin \mathcal{A}_p^-(p^*) \implies \mu_{i-}^* = 0 \quad \text{and} \quad i \in \mathcal{A}_p^-(p^*) \implies \mu_{i-}^* \geq 0 \quad (4.3a)$$

$$i \notin \mathcal{A}_p^+(p^*) \implies \mu_{i+}^* = 0 \quad \text{and} \quad i \in \mathcal{A}_p^+(p^*) \implies \mu_{i+}^* \geq 0 \quad (4.3b)$$

$$i \notin \mathcal{A}_\tau^-(\tau^*) \implies \nu_{i-}^* = 0 \quad \text{and} \quad i \in \mathcal{A}_\tau^-(\tau^*) \implies \nu_{i-}^* \geq 0 \quad (4.3c)$$

$$i \notin \mathcal{A}_\tau^+(\tau^*) \implies \nu_{i+}^* = 0 \quad \text{and} \quad i \in \mathcal{A}_\tau^+(\tau^*) \implies \nu_{i+}^* \geq 0 \quad (4.3d)$$

where \implies denotes logical implication. That is, all inequality multipliers are non-negative; however, multipliers associated with inactive constraints are zero. Take $j \in \{1, 2, \dots, n\}$. If $p_j^- = p_j^+$, then $p_j^* = p_j^- = p_j^+$. Similarly, if $\tau_j^- = \tau_j^*$, then $\tau_j^* = \tau_j^- = \tau_j^+$. We avoid these trivial cases by assuming that $p_j^- \neq p_j^+$ and $\tau_j^- \neq \tau_j^+$. Of course, if $\tau_j^+ = \infty$, then it is impossible for $\tau_j^* = \tau_j^*$.

Preference Probabilities: First, consider the requirements on the preference probabilities. Equation (4.2) requires that

$$\frac{D^* a_j(\tau_j^*) - A^* d_j(\tau_j^*)}{(D^*)^2} = \mu_{i+}^* - \mu_{i-}^*$$

There are three cases of interest.

$p_j^* \in (p_j^-, p_j^+)$: By Equations (4.3a) and (4.3b), $\mu_{j-}^* = \mu_{j+}^* = 0$. Therefore,

$$D^* a_j(\tau_j^*) = A^* d_j(\tau_j^*) \quad (4.4a)$$

$p_j^* = p_j^-$: By Equation (4.3a), $\mu_{j-}^* \geq 0$ and $\mu_{j+}^* = 0$. Therefore,

$$D^* a_j(\tau_j^*) \leq A^* d_j(\tau_j^*) \quad (4.4b)$$

$p_j^* = p_j^+$: By Equation (4.3b), $\mu_{j-}^* = 0$ and $\mu_{j+}^* \geq 0$. Therefore,

$$D^* a_j(\tau_j^*) \geq A^* d_j(\tau_j^*) \quad (4.4c)$$

For the minimum constraint to be active, the partial derivative of J at the constraint must be negative. Similarly, for the maximum constraint to be active, the partial derivative of J at the constraint must be positive. Otherwise, the partial derivative of J should be zero. Additionally, if the minimum and maximum constraints are equal, there is no restriction on the partial derivative of J at that point. All of these conditions are intuitive and can be explained graphically.

Processing Times: Next, consider the requirements on the processing times. Equation (4.2) requires that

$$\frac{D^* p_j^* a_j'(\tau_j^*) - A^* p_j^* d_j'(\tau_j^*)}{(D^*)^2} = \nu_{i+}^* - \nu_{i-}^*$$

There are three cases of interest.

$\tau_j^* \in (\tau_j^-, \tau_j^+)$: By Equations (4.3c) and (4.3d), $\nu_{j-}^* = \nu_{j+}^* = 0$. Therefore,

$$D^* p_j^* a_j'(\tau_j^*) = A^* p_j^* d_j'(\tau_j^*) \quad (4.5a)$$

$\tau_j^* = \tau_j^-$: By Equation (4.3c), $\nu_{j-}^* \geq 0$ and $\nu_{j+}^* = 0$. Therefore,

$$D^* p_j^* a_j'(\tau_j^*) \leq A^* p_j^* d_j'(\tau_j^*) \quad (4.5b)$$

$\tau_j^* = \tau_j^+$: By Equation (4.3d), $\nu_{j-}^* = 0$ and $\nu_{j+}^* \geq 0$. Therefore,

$$D^* p_j^* a_j'(\tau_j^*) \geq A^* p_j^* d_j'(\tau_j^*) \quad (4.5c)$$

Clearly, the same interpretation applies here as applied for the requirements on optimal preference probabilities.

Second-Order Necessary Conditions

Once more, assume that the point $(p^*, \tau^*) \in \mathcal{F}$ is a local maximum of the objective function and use the notation in Equation (4.1). We also use the notation

$$J_{xy}^* \triangleq \frac{\partial^2 J}{\partial x \partial y} \Big|_{(p,\tau)=(p^*,\tau^*)} \quad A_{xy}^* \triangleq \frac{\partial^2 A}{\partial x \partial y} \Big|_{(p,\tau)=(p^*,\tau^*)} \quad D_{xy}^* \triangleq \frac{\partial^2 D}{\partial x \partial y} \Big|_{(p,\tau)=(p^*,\tau^*)}$$

Again, D^* must be assumed to be nonzero. We also assume that the functions a_i and d_i are twice continuously differentiable⁴ functions for all $i \in \{1, 2, \dots, n\}$. It is necessary that there exist Lagrange multiplier vectors $m^* \in (\mathbb{R}_{\geq 0}^n)^4$ such that the first-order necessary conditions hold and

$$\delta^\top \nabla^2 L(p^*, \tau^*, m^*) \delta \geq 0 \quad \text{for all } \delta \in \mathcal{V}(p^*, \tau^*) - \{0\} \quad (4.6)$$

That is, at the point (p^*, τ^*) , the Hessian of the Lagrangian must be positive semidefinite over the set of feasible variations at that point. The Hessian $\nabla^2 L(p^*, \tau^*, m^*)$ does not depend upon the multipliers m^* , and so it is completely characterized by $J_{p_j p_k}^*$, $J_{\tau_j \tau_k}^*$, and $J_{p_j \tau_k}^*$ for all $j, k \in \{1, 2, \dots, n\}$. Therefore, take $j, k \in \{1, 2, \dots, n\}$.

Elimination of Active Preference Probability Constraints: First, assume that $j \in \mathcal{A}_p^-(p^*) \cup \mathcal{A}_p^+(p^*)$. That is, assume that an inequality constraint on the j^{th} preference probability is active (i.e., $p_j^* = p_j^-$ or $p_j^* = p_j^+$). In this case, for all $\delta \in \mathcal{V}(p^*, \tau^*)$, $\delta_j^p = 0$. Therefore, because the feasible variations along active constraint directions are zero, $J_{p_j p_k}^*$ and $J_{p_j \tau_k}^*$ will have no impact on Equation (4.6).

Elimination of Active Processing Time Constraints: Next, instead assume that $j \in \mathcal{A}_\tau^-(\tau^*) \cup \mathcal{A}_\tau^+(\tau^*)$. That is, assume that an inequality constraint on the j^{th}

⁴That is, the derivatives at each point in their domain are themselves continuously differentiable.

processing time is active (i.e., $\tau_j^* = \tau_j^-$ or $\tau_j^* = \tau_j^+$). In this case, for all $\delta \in \mathcal{V}(p^*, \tau^*)$, $\delta_j^\tau = 0$. Therefore, because the feasible variations along active constraint directions are zero, $J_{p_k \tau_j}^*$ and $J_{\tau_j \tau_k}^*$ will have no impact on Equation (4.6).

Elimination of Off-Diagonal Terms: By the reasoning about active constraints above, we can focus on coordinates of (p^*, τ^*) where constraints are inactive, and so we assume Equations (4.4a) and (4.5a). Therefore,

$$J_{p_j p_k}^* = \frac{D^* A_{p_j p_k}^* - A^* D_{p_j p_k}^*}{(D^*)^2} \quad \text{and} \quad J_{p_j \tau_k}^* = \frac{D^* A_{p_j \tau_k}^* - A^* D_{p_j \tau_k}^*}{(D^*)^2} \quad (4.7)$$

and

$$J_{\tau_j \tau_k}^* = \frac{D^* A_{\tau_j \tau_k}^* - A^* D_{\tau_j \tau_k}^*}{(D^*)^2} \quad (4.8)$$

For the moment, we focus on the off-diagonal terms of the Hessian that correspond to inactive constraints. First, assume that $j \neq k$. Clearly,

$$A_{p_j p_k}^* = D_{p_j p_k}^* = A_{\tau_j \tau_k}^* = D_{\tau_j \tau_k}^* = A_{p_j \tau_k}^* = D_{p_j \tau_k}^* = 0$$

Thus,

$$J_{p_j p_k}^* = J_{\tau_j \tau_k}^* = J_{p_j \tau_k}^* = 0$$

Now we focus on the remaining off-diagonal terms. That is, take $j = k$. So,

$$J_{p_j \tau_j}^* = \frac{D^* a_j'(\tau_j^*) - A^* d_j'(\tau_j^*)}{(D^*)^2}$$

Recall that we are taking $j \notin \mathcal{A}_p^-(p^*) \cup \mathcal{A}_p^+(p^*)$ (i.e., the j^{th} preference probability is unconstrained, so $p_j^* \in (p_j^-, p_j^+)$). Therefore, $p_j^* > 0$ and so Equation (4.5a) implies that $a_j'(\tau_j^*) = J^* d_j'(\tau_j^*)$. However, $D^* J^* = A^*$. Thus, by substitution, it is clear that $J_{p_j \tau_j}^* = 0$. Hence, $J_{p_j p_k}^*$, $J_{\tau_j \tau_k}^*$, and $J_{p_i \tau_j}^*$ have no impact on Equation (4.6) for all $i, j, k \in \{1, 2, \dots, n\}$ with $j \neq k$.

Impact of Inactive Preference Probability Diagonals: Next, we consider the diagonal terms of the Hessian that correspond to inactive preference probabilities. That is, assume that $j = k$ and $p_j^* \in (p_j^-, p_j^+)$. The condition in Equation (4.6) requires that $J_{p_j p_j}^* \leq 0$. By Equation (4.7), this means that

$$D^* \times 0 \leq A^* \times 0 \tag{4.9}$$

which is always true (i.e., it is always the case that $0 \leq 0$ with equality). Therefore, this necessary condition adds no more information than Equation (4.4a).

Definiteness from Inactive Processing Time Diagonals: By the reasoning above, the only second partial derivative that can prevent Equation (4.6) from being true is $J_{\tau_j \tau_j}^*$ where $\tau_j^* \in (\tau_j^-, \tau_j^+)$. That is, the condition in Equation (4.6) requires that $J_{\tau_j \tau_j}^* \leq 0$. By Equation (4.8), this means that

$$D^* p_j^* a_j''(\tau_j^*) \leq A^* p_j^* d_j''(\tau_j^*) \tag{4.10}$$

If the constraint parameter $p_j^- = 0$ and the j^{th} preference probability constraint is active (i.e., $p_j^* = 0$), then this condition is always true by equality. Otherwise, if $p_j^* > 0$, it must be that $D^* a_j''(\tau_j^*) \leq A^* d_j''(\tau_j^*)$.

Second-Order Sufficiency Conditions

Now take an arbitrary feasible point $(p^*, \tau^*) \in \mathcal{F}$ that may be a maximum of the objective function. If there exist Lagrange multiplier vectors $m^* \in (\mathbb{R}_{\geq 0}^n)^4$ such that Equation (4.2) holds and Equations (4.6) and (4.3a)–(4.3c) hold with *strict* inequality, then the point must be a local maximum of the objective function. This is effectively a statement of the local concavity of the objective function at the point (p^*, τ^*) .

The Extreme-Preference Rule: In order for Equation (4.6) to hold with strict inequality, Equations (4.9) and (4.10) must both hold with strict inequality. However, this is impossible for Equation (4.9). Therefore, if there is some $i \in \{1, 2, \dots, n\}$ with $p_i^* \in (p_i^-, p_i^+)$, these conditions cannot be used to show that the point is a local maximum⁵. Our goal is to design strategies guaranteed to be local maxima, so these strategies will have $p_i^* = p_i^-$ or $p_i^* = p_i^+$ for all $i \in \{1, 2, \dots, n\}$. We call this the *extreme-preference rule (EPR)*. Stephens and Krebs [60] assume that $(p_i^-, p_i^+) = (0, 1)$ for all $i \in \{1, 2, \dots, n\}$, so they call this the *zero-one rule*. This rule is part of a sufficiency condition; it is not at all necessary.

Problems with Sufficiency at Zero Preference Probability: Assume there exists some $j \in \{1, 2, \dots, n\}$ such that $p_j^* = 0$. Equations (4.5b), (4.5c), and (4.10) cannot all hold with strict inequality for this p^* . In other words, strict concavity is impossible at this point because the objective function is the same value for any choice of τ_j^* . However, it can be shown that if these all hold when p_j^* is replaced with some arbitrarily small ε with $0 < \varepsilon < p_j^+$, then the point (p_j^*, τ_j^*) is a local maximum of the objective function. In other words, even if the function is not strictly locally concave, under these ε -conditions it is certainly locally concave.

4.1.3 Solutions to Special Cases

Solutions to this generalized optimization problem can be difficult to find. In fact, mere existence of solutions cannot be taken for granted. However, there are two special cases that guarantee existence (but not uniqueness) of solutions and can be equipped with simple methods of finding one of those solutions.

⁵In other words, strict concavity cannot hold at such a point.

Constant Disadvantage Case

This case not only serves as an important example but is useful in some real cases. It is our goal to construct a strategy $(p^*, \tau^*) \in \mathcal{F}$ that meets all sufficiency conditions to be called a local maximum point of the objective function. This point will be a *global* maximum if the objective function is concave. The point will be the unique global maximum if the objective function is strictly concave. Assume that for all $j, k \in \{1, 2, \dots, n\}$,

(i) $p_j^- = 0$

(ii) a_j and d_j are twice continuously differentiable functions

(iii) for all $\tau_j \in \mathbb{R}_{\geq 0} \cap [\tau_j^-, \tau_j^+]$ and all $\tau_k \in \mathbb{R}_{\geq 0} \cap [\tau_k^-, \tau_k^+]$,

- $d(\tau_j)d \geq 0$
- $d_j(\tau_j)d_k(\tau_k) > 0$

(iv) either $d \neq 0$ or there exists some $i \in \{1, 2, \dots, n\}$ such that $p_i^* > 0$

(v) $d'_j(\tau_j) = 0$ for all $\tau_j \in (\tau_j^-, \tau_j^+)$

(vi) if $\tau_j^- \neq \tau_j^+$, it is the case that

(a) $d_j(\tau_j^-)a'_j(\tau_j^-) < 0$ or

(b) $d_j(\tau_j^+)a'_j(\tau_j^+) > 0$ or

(c) $d_j(\tau_j)a'_j(\tau_j) = 0$ with $d_j(\tau_j)a''_j(\tau_j) < 0$ for some $\tau_j \in (\tau_j^-, \tau_j^+)$

If these assumptions do not hold, for each $j \in \{1, 2, \dots, n\}$, τ_j^- and τ_j^+ may be adjusted to surround a region where they do hold. These assumptions lead to the following for all $j \in \{1, 2, \dots, n\}$.

Well-Defined Objective Function: By (iii) and (iv), $D^* \neq 0$ and $D^*d_j(\tau_j) > 0$.

This implies that both J^* and $a_j(\tau_j)/d_j(\tau_j)$ are well-defined for all choices of $\tau_j \in [\tau_j^-, \tau_j^+]$.

Maximum Type-Advantage-to-Type-Disadvantage Ratio Exists: By (vi), there exists some $\tau_j^* \in [\tau_j^-, \tau_j^+]$ such that there is some $\delta_j \in \mathbb{R}_{>0}$ where $a_j(\tau_j)/d_j(\tau_j) \leq a_j(\tau_j^*)/d_j(\tau_j^*)$ for all $\tau_j \in (\tau_j - \delta_j, \tau_j + \delta_j) \cap [\tau_j^-, \tau_j^+]$. That is, the a_j/d_j function has a maximum on its domain.

Parameterized Processing Times: If $\tau_j^- = \tau_j^+$, then (v) and (vi) are trivially met.

This case is useful when processing times are parameters of the system and not decision variables. Stephens and Krebs [60] use the name *prey model* for the case where no processing times are free decision variables (i.e., tasks are whole items of prey that come lumped with a rigid (average) processing time).

If $\tau_j^- = \tau_j^+$, let $\tau_j^* = \tau_j^-$. Otherwise, let τ_j^* be a maximum of a_j/d_j that is described by (vi). Next, assume that the types are indexed so that

$$\frac{a_1(\tau_1^*)}{d_1(\tau_1^*)} > \frac{a_2(\tau_2^*)}{d_2(\tau_2^*)} > \dots > \frac{a_{n-1}(\tau_{n-1}^*)}{d_{n-1}(\tau_{n-1}^*)} > \frac{a_n(\tau_n^*)}{d_n(\tau_n^*)} \quad (4.11)$$

Assume that for all $k \in \{0, 1, 2, \dots, n-1\}$,

$$\frac{a + \sum_{i=1}^k p_i^+ a_i(\tau_i^*)}{d + \sum_{i=1}^k p_i^+ d_i(\tau_i^*)} \neq \frac{a_{k+1}(\tau_{k+1}^*)}{d_{k+1}(\tau_{k+1}^*)}$$

Finally, define k^* by

$$k^* \triangleq \min \left(\left\{ k \in \{0, 1, 2, \dots, n-1\} : \frac{a + \sum_{i=1}^k p_i^+ a_i(\tau_i^*)}{d + \sum_{i=1}^k p_i^+ d_i(\tau_i^*)} > \frac{a_{k+1}(\tau_{k+1}^*)}{d_{k+1}(\tau_{k+1}^*)} \right\} \cup \{n\} \right)$$

and let

$$p_j^* = \begin{cases} p_j^+ & \text{if } j \leq k^* \\ 0 & \text{if } j > k^* \end{cases}$$

for all $j \in \{1, 2, \dots, n\}$. Primarily because of assumption (vi) and the results that $D^*d_j(\tau_j^*) > 0$ and $d'_j(\tau_j^*)' = d''_j(\tau_j^*) = 0$ for all $j \in \{1, 2, \dots, n\}$, it is easy to show that (p^*, τ^*) meets the conditions described in [Section 4.1.2](#) that guarantee it is a local maximum of the objective function⁶.

Decreasing Advantage-to-Disadvantage Ratio

Again, it is our goal to construct a strategy $(p^*, \tau^*) \in \mathcal{F}$ that meets all sufficiency conditions to be called a local maximum point of the objective function. However, here we assume that the disadvantage functions are not constant with respect to processing time. This is a generalized version of the *combined prey and patch model* discussed by Stephens and Krebs [60], and so it shows the MVT concept [14, 16]. However, [Stephens and Krebs](#) make different assumptions than we do because they depend on search costs being nil. Assume that for all $j, k \in \{1, 2, \dots, n\}$,

(i) $p_j^- = 0$

(ii) a_j and d_j are twice continuously differentiable functions

(iii) for all $\tau_j \in \mathbb{R}_{\geq 0} \cap [\tau_j^-, \tau_j^+]$ and all $\tau_k \in \mathbb{R}_{\geq 0} \cap [\tau_k^-, \tau_k^+]$,

- $d(\tau_j)d \geq 0$
- $d_j(\tau_j)d_k(\tau_k) > 0$

(iv) either $d \neq 0$ or there exists some $i \in \{1, 2, \dots, n\}$ such that $p_i^* > 0$

⁶Because $p_j^- = 0$ for all $j \in \{1, 2, \dots, n\}$, this statement requires the zero preference probability modification described at the end of [Section 4.1.2](#).

(v) $d_j(\tau_j)d'_j(\tau_j) > 0$ for all $\tau_j \in (\tau_j^-, \tau_j^+)$

(vi) $(a_j(\tau_j)/d_j(\tau_j))' < 0$ for all $\tau_j \in (\tau_j^-, \tau_j^+)$

(vii) $(a'_j(\tau_j)/d'_j(\tau_j))' < 0$ for all $\tau_j \in (\tau_j^-, \tau_j^+)$

If these assumptions do not hold, for each $j \in \{1, 2, \dots, n\}$, τ_j^- and τ_j^+ may be adjusted to surround a region where they do. These assumptions lead to the following for all $j \in \{1, 2, \dots, n\}$.

Well-Defined Objective Function: By (iii) and (iv), $D^* \neq 0$ and $D^*d_j(\tau_j) > 0$. This implies that both J^* , $a_j(\tau_j)/d_j(\tau_j)$, and $a'_j(\tau_j)/d'_j(\tau_j)$ are all well-defined for all choices of $\tau_j \in [\tau_j^-, \tau_j^+]$.

Maximum Type-Advantage-to-Type-Disadvantage Ratio Exists: By (vi), τ_j^- is such that $a_j(\tau_j)/d_j(\tau_j) \leq a_j(\tau_j^-)/d_j(\tau_j^-)$ for all $\tau_j \in [\tau_j^-, \tau_j^+]$. That is, the $a_j(\tau_j)/d_j(\tau_j)$ function achieves its maximum at $\tau_j = \tau_j^-$.

Ordering of Ratios: By (vi) and (v), $a_k(\tau_k)/d_k(\tau_k) > a'_k(\tau_k)/d'_k(\tau_k)$ for all $\tau_k \in (\tau_k^-, \tau_k^+)$.

Parameterized Processing Times: If $\tau_j^- = \tau_j^+$, then (v)–(vii) are trivially met.

Assume the types are indexed so that

$$\frac{a_1(\tau_1^-)}{d_1(\tau_1^-)} > \frac{a_2(\tau_2^-)}{d_2(\tau_2^-)} > \dots > \frac{a_{n-1}(\tau_{n-1}^-)}{d_{n-1}(\tau_{n-1}^-)} > \frac{a_n(\tau_n^-)}{d_n(\tau_n^-)} \quad (4.12)$$

In other words, as in the constant advantage case, order the task types by decreasing maximum advantage-to-disadvantage ratio. This is the same ordering used by [Stephens and Krebs](#); however, because we have assumed the derivative of this ratio

is strictly decreasing, the initial ratio will always be the maximum ratio. Next, for all $k \in \{0, 1, \dots, n\}$, define τ_j^k so that

$$\frac{a'_j(\tau_j^k)}{d'_j(\tau_j^k)} > \frac{a + \sum_{i=1}^k p_i^+ a_i(\tau_i^k)}{d + \sum_{i=1}^k p_i^+ d_i(\tau_i^k)} \quad \text{for } \tau_j^k = \tau_j^+$$

or

$$\frac{a'_j(\tau_j^k)}{d'_j(\tau_j^k)} < \frac{a + \sum_{i=1}^k p_i^+ a_i(\tau_i^k)}{d + \sum_{i=1}^k p_i^+ d_i(\tau_i^k)} \quad \text{for } \tau_j^k = \tau_j^-$$

or

$$\frac{a'_j(\tau_j^k)}{d'_j(\tau_j^k)} = \frac{a + \sum_{i=1}^k p_i^+ a_i(\tau_i^k)}{d + \sum_{i=1}^k p_i^+ d_i(\tau_i^k)} \quad \text{for } \tau_j^k \in (\tau_j^-, \tau_j^+)$$

By (vii), this is always possible. Unfortunately, for each $k \in \{0, 1, \dots, n\}$, all elements of the set $\{\tau_j^k : j = \{1, 2, \dots, k\}\}$ must be determined simultaneously. This is different from the constant disadvantage case. That is, because $d'_k(\tau_j) \neq 0$ for all $\tau_j \in [\tau_j^-, \tau_j^+]$, there is coupling among the optimal choices of processing time. It must also be assumed that

$$\frac{a + \sum_{i=1}^k p_i^+ a_i(\tau_i^k)}{d + \sum_{i=1}^k p_i^+ d_i(\tau_i^k)} \neq \frac{a_{k+1}(\tau_{k+1}^-)}{d_{k+1}(\tau_{k+1}^-)}$$

for all $k \in \{0, 1, 2, \dots, n-1\}$. Now, define k^* by

$$k^* \triangleq \min \left(\left\{ k \in \{0, 1, 2, \dots, n-1\} : \frac{a + \sum_{i=1}^k p_i^+ a_i(\tau_i^k)}{d + \sum_{i=1}^k p_i^+ d_i(\tau_i^k)} > \frac{a_{k+1}(\tau_{k+1}^-)}{d_{k+1}(\tau_{k+1}^-)} \right\} \cup \{n\} \right)$$

Finally, let

$$\tau_j^* = \tau_j^{k^*} \quad \text{and} \quad p_j^* = \begin{cases} p_j^+ & \text{if } j \leq k^* \\ 0 & \text{if } j > k^* \end{cases}$$

for all $j \in \{1, 2, \dots, n\}$. Primarily because of assumptions (vi) and (vii), it is easy to show that (p^*, τ^*) meets the conditions described in Section 4.1.2 that guarantee it is a local maximum of the objective function⁷.

4.2 Optimization of Specific Objective Functions

The optimization results given in Section 4.1 may be applied to many of the functions introduced in Chapter 3. We consider three of them here. Unfortunately, the reward-to-variability and reward-to-variance optimization functions do not fit the form of Section 4.1 because the central moments used to define them involve a great deal of cross-coupling among task-type parameters and decision variables. Therefore, we do not consider solutions to these optimization functions. We also do not provide solutions for the constrained optimization functions; however, we have shown other ways to implement success thresholds that can be handled by the methods in Section 4.1.

4.2.1 Maximization of Rate of Excess Net Point Gain

Consider the function $(E(G_1) - G^T/N^p)/E(T_1)$ where $G^T \in \mathbb{R}$ is a net gain success threshold. Using the statistics derived in Chapter 2, this can be expressed by

$$\frac{E(G_1) - \frac{G^T}{N^p}}{E(T_1)} = \frac{\overline{g^p} - \overline{c^p} - \frac{c^s}{\lambda^p} - \frac{G^T}{N^p}}{\overline{\tau^p} + \frac{1}{\lambda^p}} = \frac{-c^s + \sum_{i=1}^n p_i \lambda_i \left(g_i(\tau_i) - c_i \tau_i - \frac{G^T}{N^p} \right)}{1 + \sum_{i=1}^n p_i \lambda_i \tau_i}$$

Define

$$\begin{aligned} a &\triangleq -c^s & a_j(\tau_j) &\triangleq \lambda_j \left(g_j(\tau_j) - c_j \tau_j - \frac{G^T}{N^p} \right) \\ d &\triangleq 1 & d_j(\tau_j) &\triangleq \lambda_j \tau_j \end{aligned}$$

⁷Because $p_j^- = 0$ for all $j \in \{1, 2, \dots, n\}$, this statement requires the zero preference probability modification described at the end of Section 4.1.2.

Using these definitions, $(E(G_1) - G^T/N^p)/E(T_1)$ fits the form studied in [Section 4.1](#).

4.2.2 Maximization of Discounted Net Gain

Consider the function $E(G_1) - w E(T_1)$ where $w \in \mathbb{R}$. Using the statistics derived in [Chapter 2](#), this can be expressed by

$$E(G_1) - w E(T_1) = \overline{g^p} - \overline{c^p} - \frac{c^s}{\lambda^p} - w \overline{\tau^p} - w \frac{1}{\lambda^p} = \frac{-(c^s + w) + \sum_{i=1}^n p_i \lambda_i (g_i(\tau_i) - c_i \tau_i - w \tau_i)}{\sum_{i=1}^n p_i \lambda_i}$$

Define

$$\begin{aligned} a &\triangleq -(c^s + w) & a_j(\tau_j) &\triangleq \lambda_j (g_j(\tau_j) - c_j \tau_j - w \tau_j) \\ d &\triangleq 0 & d_j(\tau_j) &\triangleq \lambda_j \end{aligned}$$

Using these definitions, clearly $E(G_1) - w E(T_1)$ fits the form studied in [Section 4.1](#). This is a constant disadvantage example. Consider fixing processing times to be parameters so that the excess rate of gain function in [Section 4.2.1](#) is also a constant disadvantage example. Also take $G^T = 0$. In this case, the resulting rate of net gain function is nearly identical to the one studied in classical OFT. In this constant disadvantage context (called the *prey model* by Stephens and Krebs [60]), indexing by advantage-to-disadvantage ratio will often lead to the same ordering for both the rate of net gain and the discounted net gain functions. Therefore, if observational justification for the use of rate of point gain as an optimization objective is based entirely on task-type ranking, then discounted net gain is an equally valid optimization objective to consider.

4.2.3 Maximization of Rate of Excess Efficiency

Consider the function $(E(G_1) + E(C_1) - G_g^T/N^p)/E(C_1)$ where $G_g^T \in \mathbb{R}$ is a gross gain success threshold. Using the statistics derived in [Chapter 2](#), this can be expressed by

$$\frac{E(G_1) + E(C_1) - \frac{G_g^T}{N^p}}{E(C_1)} = \frac{\bar{g}^p - \frac{G_g^T}{N^p}}{\bar{c}^p + \frac{c^s}{\lambda^p}} = \frac{\sum_{i=1}^n p_i \lambda_i \left(g_i(\tau_i) - \frac{G_g^T}{N^p} \right)}{c^s + \sum_{i=1}^n p_i \lambda_i c_i \tau_i}$$

Define

$$\begin{aligned} a &\triangleq 0 & a_j(\tau_j) &\triangleq \lambda_j \left(g_j(\tau_j) - \frac{G_g^T}{N^p} \right) \\ d &\triangleq c^s & d_j(\tau_j) &\triangleq \lambda_j c_j \tau_j \end{aligned}$$

Using these definitions, $(E(G_1) + E(C_1) - G_g^T/N^p)/E(C_1)$ fits the form studied in [Section 4.1](#).

There are two major criticisms of optimizing efficiency [[60](#), p. 9]. First, it ignores the impact of time. Second, it equates behaviors that bring small gains at small costs with behaviors that bring large gains at large costs. Together, an efficiency optimizer can spend large amounts of time for a small gain that is insufficient for survival. However, costs in our model are affinely related to time, so cost minimization exerts pressure on time as well. Additionally, efficiency is defined with a success threshold (i.e., excess efficiency), and so all behaviors that have positive efficiency also lead to survival. Therefore, if our model can be used, efficiency maximization may be a viable alternative to rate maximization. In a constant disadvantage context, the efficiency advantage-to-disadvantage indexing will be very similar to the indexing in [Section 4.2.1](#), and so evidence for the use of rate maximization may also justify the use of efficiency maximization.

CHAPTER 5

CONCLUSION

With increasing demand for automation, engineering design methods must be developed that encapsulate complex high-level decision making. Automated controllers need to perform tasks that would traditionally be carried out by cognitive agents (e.g., human beings). So, engineering is progressively more interested in constructing *behaviors* rather than just decision rules. Therefore, it makes sense that behavioral ecology could be influential to the development of design methods in artificial intelligence. This insight is the genesis of this thesis. We demonstrate how the study of patterns of natural cognitive behavior can be used to guide the construction of engineered agents. This novel extension of behavioral ecology can lead to new realizations about the natural world. As these fields have a rich history, their combination provides for many future research directions.

5.1 Contributions to Engineering

Results from behavioral ecology can be extended to engineering design. At an abstract level, a forager with a behavior favored by natural selection due to its energy-time balance is no different from a single agent with protocols that achieve a favorable

work-resource balance. This agent model is applicable in a wide range of engineering applications. The analogy between foraging and task processing is obvious in military applications where agents search for targets to process while also minimizing fuel cost or risk. Also consider an automated centralized temperature controller in a large building. The controller has limited control authority and faces random temperature disturbances. It must prioritize its efforts to achieve some desirable temperature profile given its limited resources. Our design methods may be used to design a strategy that efficiently manages the temperature profile of the building¹. In fact, Quijano et al. [50] have implemented an OFT-based temperature controller prototype. Other controllers that must prioritize resource investment to achieve some favorable outcome may be viewed in a similar way. Our model is particularly useful in applications with Poisson encounters, as in many queueing applications. This model could also be modified for use in other stochastic environments.

5.2 Contributions to Biology

This new engineering approach is not only inspired by classical OFT but also provides new insights to behavioral ecology. While qualitative results can be useful in justifying observed behavior in nature, the design of engineered behaviors requires a strong quantitative analysis of a trusted model. Therefore, the two fields have different priorities, and they complement each other. Biologically-inspired engineering design leads to elegant agent behaviors and new insights into the elegance of observed foraging behaviors in nature. In particular, the work in this thesis contributes to biology in the following ways.

¹For example, temperature perturbations away from the desired profile may be viewed as task encounters.

Improved Agent Model: Our solitary agent model that we discuss encapsulates the existing foraging model used in classical OFT. However, we explicitly model a relationship between processing time and processing cost and none of our analysis requires that any cost is nil. We also allow for the possibility of negative search costs and gains, which expands the applicability of the model (e.g., negative search costs may indicate additional value accumulated while not in a processing mode).

Combination of Rate Maximization and Risk Sensitivity: Our approach of defining an agent lifetime in terms of a finite number of tasks yields new ways of approaching statistical optimization of the agent model. Because the agent has a finite lifetime, thresholds of lifetime success may be added to the analysis. Success thresholds cannot be used in classical OFT rate maximization because an infinite lifetime is assumed, and so any threshold will have zero impact on behavior. Gain thresholds are considered in risk-sensitive classical OFT approach, but the impact of environmental parameters cannot easily be explored because the model does not provide an easy way to study finite-lifetime behavior. Because our approach is defined by a finite lifetime assumption, we can add a gain threshold to rate maximization (i.e., we consider *excess* rate maximization) and study the impact of environmental parameter changes on risk-sensitive behavior. The former combines risk sensitivity and time minimization. The latter shows not only how parameters like encounter rates and costs can modulate risk-sensitive behavior but also how minimization of uncertainty is related to time minimization.

Evolutionary Justification for Efficiency Maximization: Efficiency maximization is usually considered to be unrealistic in biology because it does not provide the time minimization important to both risk sensitivity and rate maximization. However, by studying efficiency with respect to our agent model, we show how its maximization does provide time pressure (i.e., minimization of costs has a related effect on time). This makes it an optimization objective that does not conflict with the expected pressures of natural selection or survival in general.

Generalized MVT: By studying the maximization of a generalized rational value function, we provide a simple method for finding behaviors that are optimal with respect to existing objectives and objectives yet to be determined. The optimal solutions to this generic value function show that the MVT is a specialization of a general rule based on marginal advantage and marginal disadvantage.

These contributions are the result of a fresh perspective on well-known theoretical research in behavioral ecology. This suggests that collaboration between engineers and biologists has synergistic value.

5.3 Future Directions

There are several future directions for extending this work. For one, the agent model we have described may be expanded to include the impact of recognition cost and behavior-dependent encounter rates. Nonlinear fuel costs might also be added to the model². Additionally, analytical results that use variance would be valuable when considering risk in random environments. The optimization methods also leave

²As our justifications for using efficiency maximization as an alternative for rate maximization are based on linear costs, this modification is not trivial.

room for improvement. As discussed, MPT and PMPT have studied nearly identical problems in finance. Modern portfolio choice and capital budgeting research is far more advanced than the economic literature typically cited by behavioral ecologists. Approaching behavioral analysis and design from this updated point of view may be valuable. Finally, it is important to engineering that biologically-inspired agent design be tested experimentally in order to validate its utility.

5.4 The Value of Collaboration

Studying ways of combining behavioral ecology, finance, and engineering has been enlightening and stimulating. Researchers in these fields approach similar problems from different directions. Their collaboration can lead to unanticipated insights of genuine value. Even if the results of this particular work fail to be successfully applied to engineering, it is likely that starting a discussion among members of this diverse group of fields will eventually yield mutual benefits.

APPENDIX A

LIMITS OF MARKOV RENEWAL PROCESSES

Take a probability space $(\mathcal{U}, \Sigma, \Pr)$. Let $(M(t_s) : t_s \in \mathbb{R}_{\geq 0})$ be a Poisson process from the space with rate $\lambda \in \mathbb{R}_{>0}$ and interevent time process (Υ_M) . Take (τ_M) to be a sequence of non-negative random variables from this space where τ_M is independent of Υ_N for all $M, N \in \mathbb{N}$. From these, define the random processes (T_N) and (T^N) with

$$T_N \triangleq \Upsilon_N + \tau_N \quad \text{and} \quad T^N \triangleq \sum_{i=1}^N T_N = \sum_{i=1}^N \Upsilon_i + \tau_i$$

for all $N \in \mathbb{N}$ and $\zeta \in \mathcal{U}$. Define the Markov renewal process $(N(t) : t \in \mathbb{R}_{\geq 0})$ by

$$N(t) \triangleq \sup \left\{ N \in \mathbb{N} : \sum_{i=1}^N T_i \leq t \right\}$$

For all outcomes $\zeta \in \mathcal{U}$,

$$T^2 = T_1 + T_1 = \Upsilon_1 + \Upsilon_2 + \tau_1 + \tau_2 = A^2 + B^2$$

where random variables $A^2 : \mathcal{U} \mapsto \overline{\mathbb{R}}$ and $B^2 : \mathcal{U} \mapsto \overline{\mathbb{R}}$ are defined by $A^2 \triangleq \Upsilon_1 + \Upsilon_2$ and $B^2 \triangleq \tau_1 + \tau_2$ for all $\zeta \in \mathcal{U}$. For brevity, we assume that the probability measures associated with these two random variables are absolutely continuous with respect to the Lebesgue measure¹. The main result of this appendix holds for the general case

¹By the Poisson assumption, this must be true for A^2 .

as well. Because T^2 is the sum of two independent random variables, for all $x \in \overline{\mathbb{R}}$, $f_{T^2}(x) = f_{A^2}(x) * f_{B^2}(x)$ where the $*$ operator denotes convolution [45]. Therefore,

$$\begin{aligned} \mathbb{E}\left(\frac{1}{T^2}\right) &= \int_{-\infty}^{\infty} \frac{1}{x} (f_{A^2}(x) * f_{B^2}(x)) \, dx \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{1}{x} f_{A^2}(x-t) f_{B^2}(t) \, dt \, dx \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{1}{x} f_{A^2}(x-t) f_{B^2}(t) \, dx \, dt \\ &= \int_{-\infty}^{\infty} f_{B^2}(t) \int_{-\infty}^{\infty} \frac{1}{y+t} f_{A^2}(y) \, dy \, dt \end{aligned}$$

However, f_{B^2} is a probability density where $f_{B^2}(t) = 0$ for all $t < 0$, and so

$$\mathbb{E}\left(\frac{1}{T^2}\right) \leq \int_{-\infty}^{\infty} \frac{1}{y} f_{A^2}(y) \, dy$$

Because (Υ_M) is the interevent time process for Poisson process $(M(t_s) : t_s \in \mathbb{R}_{\geq 0})$, A^2 is Erlang-2 distributed² with parameter λ . Thus, $\mathbb{E}(1/T^2) \leq \lambda$. So, there exists an $N \in \mathbb{N}$ such that $\mathbb{E}(1/T^N) < \infty$. By results of Johns and Miller [28], for all $K \in \mathbb{N}$,

$$\text{aslim}_{t \rightarrow \infty} \frac{N(t)}{t} = \lim_{t \rightarrow \infty} \frac{\mathbb{E}(N(t))}{t} = \text{aslim}_{N \rightarrow \infty} \frac{N}{T^N} = \lim_{N \rightarrow \infty} \mathbb{E}\left(\frac{N}{T^N}\right) = \frac{K}{\mathbb{E}(T^K)} = \frac{1}{\mathbb{E}(T_1)}$$

where

$$\frac{1}{\mathbb{E}(T_1)} = \frac{1}{\mathbb{E}(\Upsilon_1) + \mathbb{E}(\tau_1)} = \frac{1}{\frac{1}{\lambda} + \mathbb{E}(\tau_1)}$$

This could be called the *long-term encounter rate* of $(N(t) : t \in \mathbb{R}_{\geq 0})$. Now define the process $(T(t) : t \in \mathbb{R}_{\geq 0})$ by

$$T(t) \triangleq T^{N(t)} = \sum_{i=1}^{N(t)} T_1$$

for all $t \in \mathbb{R}_{\geq 0}$. It is similarly the case that

$$\text{aslim}_{t \rightarrow \infty} \frac{T(t)}{t} = \lim_{t \rightarrow \infty} \frac{\mathbb{E}(T(t))}{t} = 1$$

²The Erlang-2 distribution is characterized by density $f_{A^2}(y) = \lambda^2 y \exp(-\lambda y)$ for $y > 0$.

BIBLIOGRAPHY

- [1] B. W. Andrews, K. M. Passino, and T. A. Waite. Foraging theory for decision-making system design: task-type choice. In *Proceedings, 43rd IEEE Conference on Decision and Control*, volume 5, pages 4740–4745, December 2004. <1>
- [2] B. W. Andrews, K. M. Passino, and T. A. Waite. Social foraging theory for robust multiagent system design. *IEEE Transactions on Automation Science and Engineering*, 4(1):79–86, January 2007. <1>
- [3] H. R. Arkes and P. Ayton. The sunk cost and concorde effects: Are humans less rational than lower animals? *Psychological Bulletin*, 125(5):591–600, 1999. <55>
- [4] M. Bateson and A. Kacelnik. Preferences for fixed and variable food sources: Variability in amount and delay. *Journal of the Experimental Analysis of Behavior*, 63(3):313–329, May 1995. <49, 50>
- [5] M. Bateson and S. C. Whitehead. The energetic costs of alternative rate currencies in the foraging starling. *Ecology*, 77(4):1303–1307, June 1996. <49>
- [6] V. S. Bawa. Optimal rules for ordering uncertain prospects. *Journal of Financial Economics*, 2:95–121, 1975. <45, 70, 72, 73>
- [7] V. S. Bawa. Safety-first, stochastic dominance, and optimal portfolio choice. *The Journal of Financial and Quantitative Analysis*, 13(2):255–271, June 1978. <71>
- [8] V. S. Bawa. Stochastic dominance: A research bibliography. *Management Science*, 28(6):698–712, June 1982. <73>
- [9] V. S. Bawa and E. B. Lindenberg. Capital market equilibrium in a mean-lower partial moment framework. *Journal of Financial Economics*, 5:189–200, 1977. <71>
- [10] D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, Belmont, Mass., 1995. <36, 67, 77>

- [11] T. Caraco. On foraging time allocation in a stochastic environment. *Ecology*, 61(1):119–128, 1980. <66, 73>
- [12] Y. Carmel and Y. Ben-Haim. Info-gap robust-satisficing model of foraging behavior: Do foragers optimize or satisfice? *The American Naturalist*, 166(5):633–641, November 2005. <49>
- [13] G. Chamberlain. A characterization of the distributions that imply mean-variance utility functions. *Journal of Economic Theory*, 29:185–201, 1983. <45>
- [14] E. L. Charnov. Optimal foraging: The marginal value theorem. *Theoretical Population Biology*, 9(2):129–136, April 1976. <3, 44, 87>
- [15] E. L. Charnov. Optimal foraging: Attack strategy of a mantid. *The American Naturalist*, 110(971):141–151, January–February 1976. <3>
- [16] E. L. Charnov and G. H. Orians. *Optimal foraging: some theoretical explorations*. PhD thesis, University of Washington, 1973. <3, 5, 13, 42, 44, 50, 87>
- [17] P. Chunchindaa, K. Dandapanib, S. Hamidb, and A. J. Prakash. Portfolio selection and skewness: Evidence from international stock markets. *Journal of Banking & Finance*, 21:143–167, 1997. <34>
- [18] C. W. Clark and M. Mangel. *Dynamic State Variable Models in Ecology*. Oxford University Press, New York, 2000. <50>
- [19] M. L. Cody. Optimization in ecology. *Science*, 183(4130):1156–1164, March 1974. <31>
- [20] S. Engen and N. C. Stenseth. A general version of optimal foraging theory: the effect of simultaneous encounters. *Theoretical Population Biology*, 26(2):192–204, 1984. <41>
- [21] R. P. Gendron and J. E. R. Staddon. Searching for cryptic prey: The effect of search rate. *The American Naturalist*, 121(2):172–186, February 1983. <37>
- [22] H. Grootveld and W. Hallerbach. Variance vs downside risk: Is there really that much difference? *European Journal of Operational Research*, 114:304–319, 1999. <71>
- [23] L. D. Harder and L. A. Real. Why are bumble bees risk averse? *Ecology*, 68(4):1104–1108, 1987. <49, 50>
- [24] A. I. Houston and J. M. McNamara. *Models of Adaptive Behavior*. Cambridge Univ. Press, Cambridge, 1999. <31, 41, 42, 50, 64>

- [25] R. N. Hughes. Optimal diets under the energy maximization premise: The effects of recognition time and learning. *The American Naturalist*, 113(2):209–221, February 1979. <3>
- [26] J. M. C. Hutchinson and G. Gigerenzer. Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behavioural Processes*, 69:97–124, 2005. <49>
- [27] Y. Iwasa, M. Higashi, and N. Yamamura. Prey distribution as a factor determining the choice of optimal foraging strategy. *The American Naturalist*, 117(5):710–723, May 1981. <49>
- [28] M. V. Johns and R. G. Miller, Jr. Average renewal loss rates. *The Annals of Mathematical Statistics*, 34(2):396–401, June 1963. <19, 99>
- [29] A. Kane. Skewness preference and portfolio choice. *Journal of Financial and Quantitative Analysis*, 17(1):15–25, March 1982. <34>
- [30] H. Konno, H. Shirakawa, and H. Yamazaki. A mean-absolute deviation-skewness portfolio optimization model. *Annals of Operations Research*, 45(1):205–220, December 1993. <34>
- [31] T.-Y. Lai. Portfolio selection with skewness: A multiple-objective approach. *Review of Quantitative Finance and Accounting*, 1(3):293–305, July 1991. <34>
- [32] W. Y. Lee and R. K. S. Rao. Mean lower partial moment valuation and log-normally distributed returns. *Management Science*, 34(4):446–453, April 1988. <71>
- [33] J. Lintner. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, 47(1):13–37, February 1965. <70>
- [34] M. Mangel and C. W. Clark. *Dynamic Modeling in Behavioral Ecology*. Princeton University Press, New Jersey, 1988. <50>
- [35] J. C. T. Mao. Models of capital budgeting, E-V vs E-S. *The Journal of Financial and Quantitative Analysis*, 4(5):657–675, January 1970. <71>
- [36] H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, March 1952. <1, 61, 65, 70>
- [37] H. M. Markowitz. *Portfolio Selection: Efficient Diversification of Investments*. Wiley, New York, 1959. <70>

- [38] J. Marschak. Neumann's and Morgenstern's new approach to static economics. *The Journal of Political Economy*, 54(2):97–115, April 1946. <66>
- [39] J. M. McNamara. Optimal control of the diffusion coefficient of a simple diffusion process. *Mathematics of Operations Research*, 8(3):373–380, August 1983. <47>
- [40] J. Meyer. Two-moment decision models and expected utility maximization. *The American Economic Review*, 77(3):421–430, June 1987. <45>
- [41] J. Mossin. Equilibrium in a capital asset market. *Econometrica*, 34(4):768–783, October 1966. <70>
- [42] W. W. Murdoch. Switching in general predators: Experiments on predator specificity and stability of prey populations. *Ecological Monographs*, 39(4):335–354, Autumn 1969. <37>
- [43] P. Nonacs. State dependent behavior and the marginal value theorem. *Behavioral Ecology*, 12(1):71–83, January 2001. <48, 50>
- [44] J. Owen and R. Rabinovitch. On the class of elliptical distributions and their applications to the theory of portfolio choice. *The Journal of Finance*, 38(3):745–752, June 1983. <45>
- [45] A. Papoulis and S. U. Pillai. *Probability, Random Variables and Stochastic Processes*. McGraw-Hill, New York, fourth edition, 2002. <99>
- [46] T. P. Pavlic and K. M. Passino. Foraging theory for mobile agent speed choice. *Engineering Applications of Artificial Intelligence*. Submitted. <1, 10, 37>
- [47] H. R. Pulliam. On the theory of optimal diets. *The American Naturalist*, 108(959):59–74, January–February 1974. <3>
- [48] H. R. Pulliam. Diet optimization with nutrient constraints. *The American Naturalist*, 109(970):765–768, November–December 1975. <3, 37, 38>
- [49] G. H. Pyke, H. R. Pulliam, and E. L. Charnov. Optimal foraging: A selective review of theory and tests. *The Quarterly Review of Biology*, 52(2):137–154, 1977. <38, 39, 40, 50>
- [50] N. Quijano, B. W. Andrews, and K. M. Passino. Foraging theory for multizone temperature control. *IEEE Computational Intelligence Magazine*. To appear. <1, 94>
- [51] D. J. Rapport. An optimization model of food selection. *The American Naturalist*, 105(946):575–587, November–December 1971. <36, 37, 38, 39, 53>

- [52] L. A. Real. Fitness, uncertainty, and the role of diversification in evolution and behavior. *The American Naturalist*, 115(5):623–638, May 1980. <49, 65, 66, 73>
- [53] B. M. Rom and K. W. Ferguson. Post-modern portfolio theory comes of age. *Journal of Investing*, pages 11–17, Fall 1994. <70, 71>
- [54] M. Rothschild and J. E. Stiglitz. Increasing risk: 1. A definition. *Journal of Economic Theory*, 2:225–243, September 1970. <45>
- [55] T. W. Schoener. Theory of feeding strategies. *Annual Review of Ecology and Systematics*, 2:369–404, 1971. <3, 31, 38, 39, 40, 50, 64>
- [56] W. F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3):425–442, September 1964. <70>
- [57] W. F. Sharpe. Mutual fund performance. *Journal of Business*, 39(1):119–138, January 1966. <46, 58, 61>
- [58] W. F. Sharpe. The sharpe ratio. *Journal of Portfolio Management*, 21(1):49–58, Fall 1994. <47>
- [59] D. W. Stephens and E. L. Charnov. Optimal foraging: Some simple stochastic models. *Behavioral Ecology and Sociobiology*, 10:251–263, 1982. <39, 41, 45, 47, 49>
- [60] D. W. Stephens and J. R. Krebs. *Foraging Theory*. Princeton Univ. Press, Princeton, NJ, 1986. <1, 5, 11, 12, 13, 41, 43, 46, 47, 50, 66, 73, 84, 86, 87, 88, 91, 92>
- [61] A. R. Templeton and L. R. Lawlor. The fallacy of the averages in ecological optimization theory. *The American Naturalist*, 117(3):390–393, March 1981. <49>
- [62] J. Tobin. Liquidity preference as behavior towards risk. *The Review of Economic Studies*, 25(2):65–86, February 1958. <61, 65, 70>
- [63] S. C. Tsiang. The rationale of the mean-standard deviation analysis, skewness preference, and the demand for money. *The American Economic Review*, 62(3):354–371, June 1972. <34>
- [64] Y. Viniotis. *Probability and Random Processes for Electrical Engineering*. McGraw-Hill, New York, 1998. <14>
- [65] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944. <66, 72>

- [66] E. Wajnberg, P. Bernhard, F. Hamelin, and G. Boivin. Optimal patch time allocation for time-limited foragers. *Behavioral Ecology and Sociobiology*, 60: 1–10, 2006. <51>
- [67] E. E. Werner and D. J. Hall. Optimal foraging and the size selection of prey by the bluegill sunfish (*Lepomis macrochirus*). *Ecology*, 55(5):1042–1052, 1974. <3>

LIST OF ACRONYMS¹

CLT	central limit theorem	<46>
EPR	extreme-preference rule	<84>
FSD	first-order stochastic dominance	<73>
i.i.d.	(mutually) independent and identically distributed	<8>
LPM	lower-partial moment	<71>
LPV	lower partial variance	<70>
MLPM	mean-lower-partial-moment	<71>
MLPV	mean-lower-partial-variance	<71>
MSA	mean-semivariance analysis	<71>
MVA	mean-variance analysis	<70>
MVT	marginal value theorem	<44>
OFT	optimal foraging theory	<1>
SD	stochastic dominance	<72>
TSD	third-order stochastic dominance	<73>

¹The page where the glossary entry is defined is given in angle brackets (e.g., <106>).

LIST OF TERMS¹

General Environment and Task-Type Terms

- c^s Average cost rate for searching (points per second) <8>
- n Number of task types <8>
- λ_i Poisson encounter rate for task type i (tasks per second) <10>
- c_i Average fuel cost rate for task type i (points per second per task) <10>
- g_i Average gross processing gain for task type i (points per task) <10>
- τ_i Average processing time for task type i (seconds per task) <10>
- p_i The agent's preference probability for task type i <10>
- $g_i(\tau_i)$ Average processing gain for task type i as function of average processing time <11>

Classical OFT Terms

- λ Merged Poisson encounter rate for all tasks (tasks per second) <14>
- $N(t)$ Number of tasks encountered after t seconds <16>
- g Gross gain random variable for processing a task during one OFT Markov renewal cycle (points) <14>
- \bar{g} Expected gross gain for processing a task during one OFT Markov renewal cycle (points) <14>
- c Cost random variable for processing a task during one OFT Markov renewal cycle (points) <14>
- \bar{c} Expected cost for processing a task during one OFT Markov renewal cycle (points) <14>
- τ Time random variable for processing a task during one OFT Markov renewal cycle (seconds) <14>
- $\bar{\tau}$ Expected time for processing a task during one OFT Markov renewal cycle (seconds) <14>
- \tilde{G}_1 Net gain from a single OFT renewal cycle (points) <15>
- \tilde{C}_1 Cost from a single OFT renewal cycle (points) <15>
- \tilde{T}_1 Length of time of a single OFT renewal cycle (seconds) <15>
- \tilde{G}^N Total net gain for N OFT Markov renewal cycles (points) <18>
- \tilde{C}^N Total cost for N OFT Markov renewal cycles (points) <18>

¹The page where the glossary entry is defined is given in angle brackets (e.g., <107>).

- \tilde{T}^N Total length of time for N OFT Markov renewal cycles (seconds) <17>
- $\tilde{G}(t)$ Total net gain after t seconds (i.e., $\tilde{G}^{N(t)}$) (points) <18>
- $\tilde{C}(t)$ Total cost after t seconds (i.e., $\tilde{C}^{N(t)}$) (points) <18>
- $\tilde{T}(t)$ Total length of time after t seconds for all completed OFT Markov renewal cycles (i.e., $\tilde{T}^{N(t)}$) (seconds) <18>

Processing-Only Terms

- λ_i^p Poisson encounter rate with processed tasks of type i (processed tasks per second) <22>
- λ^p Poisson encounter rate with processed tasks of all types (processed tasks per second) <23>
- $N^p(t)$ Number of tasks processed after t seconds <25>
- g^p Gross gain random variable for processing a task during one processing Markov renewal cycle (points) <23>
- $\overline{g^p}$ Expected gross gain for processing a task during one processing Markov renewal cycle (points) <24>
- c^p Cost random variable for processing a task during one processing Markov renewal cycle (points) <23>
- $\overline{c^p}$ Expected cost for processing a task during one processing Markov renewal cycle (points) <24>
- τ^p Time random variable for processing a task during one processing Markov renewal cycle (seconds) <23>
- $\overline{\tau^p}$ Expected time for processing a task during one processing Markov renewal cycle (seconds) <24>
- G_1 Net gain from a single processing renewal cycle (points) <24>
- C_1 Cost from a single processing renewal cycle (points) <24>
- T_1 Length of time of a single processing renewal cycle (seconds) <24>
- G^{N^p} Total net gain for N processing Markov renewal cycles (points) <26>
- C^{N^p} Total cost for N processing Markov renewal cycles (points) <26>
- T^{N^p} Total length of time for N^p processing Markov renewal cycles (seconds) <26>
- $G(t)$ Total net gain after t seconds (i.e., $G^{N^p(t)}$) (points) <26>
- $C(t)$ Total cost after t seconds (i.e., $C^{N^p(t)}$) (points) <26>
- $T(t)$ Total length of time after t seconds for all completed processing Markov renewal cycles (i.e., $T^{N^p(t)}$) (seconds) <26>

LIST OF SYMBOLS

Document Conventions

[**xx**] see reference number **xx** in [the bibliography](#)

General Mathematics

$=$ is equal to
 \triangleq defined as
 \approx is approximately equal to
 $<$ ($>$) strictly less (greater) than
 \leq (\geq) less (greater) than or equal to
 $x + y$ sum of x and y
 $x \times y$ product of x and y (also denoted xy)
 $-x$ additive inverse of x
 $x - y$ difference of x and y (i.e., $x - y \triangleq x + -y$)
 $\text{sgn}(x)$ sign function of x
 \prod product of elements of a set
 \sum sum of elements of a set

Numbers

\mathbb{N} the set of the natural numbers (i.e., $\{1, 2, 3, \dots\}$)
 \mathbb{W} the set of the whole numbers (i.e., $\{0, 1, 2, 3, \dots\}$)
 \mathbb{Z} the set of the integers (i.e., $\{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}$)
 \mathbb{Q} the set of the rationals (i.e., ratios of integers)
 \mathbb{R} the set of the real numbers
 $\mathbb{R}_{>0}$ the set of the strictly positive real numbers
 $\mathbb{R}_{\geq 0}$ the set of the non-negative real numbers
 $\mathbb{R}_{<0}$ the set of the strictly negative real numbers
 $\mathbb{R}_{\leq 0}$ the set of the non-positive real numbers
 $\mathbb{R}_{\neq 0}$ the set of the non-zero real numbers
 $\overline{\mathbb{R}}$ the set of the extended real numbers (i.e., $\mathbb{R} \cup \{-\infty, +\infty\}$)
 \mathbb{R}^n the Euclidean n -space
 $\mathbb{R}^{n \times m}$ space of n -by- m real matrices
 $\mathbb{R}^{n \times n}$ the unitary associative real algebra
 e Euler's number (i.e., constant $e \approx 2.71828182845904523536$)

- $\log_b(x)$ logarithm of positive real number x in base b (i.e., $b^{\log_b(x)} = x$)
- $\log(x)$ common logarithm of positive real number x (i.e., $10^{\log(x)} = x$)
- $\ln(x)$ natural logarithm of positive real number x (i.e., $e^{\ln(x)} = x$)
- $\exp(x)$ exponential function (i.e., $\exp(x) \triangleq e^x$)
- $\lceil x \rceil$ the ceiling of real number x (i.e., the least integer not less than x)
- $\lfloor x \rfloor$ the floor of real number x (i.e., the greatest integer not greater than x)

Sets

- \mathcal{X} a set \mathcal{X}
- $\{a, b, c\}$ a set of objects a , b , and c
- \dots continue the established pattern ad infinitum (e.g., the infinite set $\{1, 2, 3, \dots\}$)
- $\{u : p\}$ set of all elements of u such that p
- $\{u : p, q, r\}$ set of all elements of u such that p , q , and r
- \emptyset the empty set (i.e., $\{\}$)
- \in is an element of (i.e., set inclusion)
- \notin is not an element of (i.e., set exclusion)
- \subset (\supset) is a proper/strict subset (superset) of
- \subseteq (\supseteq) is a subset (superset) of
- $\mathcal{X} = \mathcal{Y}$ set \mathcal{X} is equal to set \mathcal{Y} (i.e., $\mathcal{X} \subseteq \mathcal{Y}$ and $\mathcal{Y} \subseteq \mathcal{X}$)
- $\mathcal{X} \neq \mathcal{Y}$ set \mathcal{X} is not equal to set \mathcal{Y}
- $\mathcal{P}(\mathcal{U})$ power set of set \mathcal{U} (i.e., the set of all subsets of \mathcal{U})
- $|\mathcal{X}|$ cardinality of set \mathcal{X}
- \cap intersection of many sets (compare to \sum)
- \cup union of many sets (compare to \sum)
- $\mathcal{X} \cap \mathcal{Y}$ set intersection (or meet) of sets \mathcal{X} and \mathcal{Y}
- $\mathcal{X} \cup \mathcal{Y}$ set union (or join) of sets \mathcal{X} and \mathcal{Y}
- $\mathcal{X} - \mathcal{Y}$ difference of sets \mathcal{X} and \mathcal{Y}
- \mathcal{X}^c complement of set \mathcal{X}^c (e.g., $\mathcal{U} - \mathcal{X}$ where $\mathcal{X} \subseteq \mathcal{U}$)
- (a, b) ordered pair of objects a and b (i.e., $(a, b) \triangleq \{\{a\}, \{a, b\}\}$)
- (x_1, x_2, \dots, x_n) n -tuple (i.e., tuple of length $n \in \mathbb{N}$ with coordinates x_1, x_2, \dots, x_n in their respective order)
- $\mathcal{X} \times \mathcal{Y}$ (binary) Cartesian product of sets \mathcal{X} and \mathcal{Y} (i.e., $\mathcal{X} \times \mathcal{Y} \triangleq \{(x, y) : x \in \mathcal{X}, y \in \mathcal{Y}\}$)
- $\mathcal{X}_1 \times \dots \times \mathcal{X}_n$ Cartesian product of n sets $\mathcal{X}_1, \dots, \mathcal{X}_n$ (i.e., $\mathcal{X}_1 \times \dots \times \mathcal{X}_n \triangleq \{(x_1, \dots, x_n) : x_1 \in \mathcal{X}_1, \dots, x_n \in \mathcal{X}_n\}$)
- \mathcal{X}^n Cartesian product of set \mathcal{X} with itself n times (e.g., $\mathcal{X}^3 \triangleq \mathcal{X} \times \mathcal{X} \times \mathcal{X}$)
- $[a, b]$ interval $[a, b] \triangleq \{x \in \mathcal{X} : a \leq x \leq b\}$
- $(a, b]$ interval $(a, b] \triangleq \{x \in \mathcal{X} : a < x \leq b\}$
- $[a, b)$ interval $[a, b) \triangleq \{x \in \mathcal{X} : a \leq x < b\}$

(a, b) interval $(a, b) \triangleq \{x \in \mathcal{X} : a < x < b\}$

Families and Sequences

- $x(i)$ or x_i or x^i alternate notations for an index i on a symbol x
 $(x_i : i \in \mathcal{I})$ an indexed family with index set \mathcal{I} (also $(x_i)_{i \in \mathcal{I}}$)
 $(x(t) : t \geq 0)$ an ordered indexed family with a directed index set \mathcal{T} where $0 \in \mathcal{T}$
 (x_α) a net (i.e., an ordered indexed family $(x_\alpha : \alpha \in \mathcal{A})$ with directed index set \mathcal{A})
 (x_n) a sequence (i.e., an ordered indexed family $(x_n : n \in \mathbb{N})$ with totally ordered index set \mathbb{N})

Logic

- \iff logical equivalence
 \implies logical implication

Order

- \inf infimum (i.e., greatest lower bound or meet)
 \sup supremum (i.e., lowest upper bound or join)
 \max maximum element
 \min minimum element

Functions and Real Analysis

- $n!$ factorial of n (i.e., $n! = 1 \times 2 \times \cdots \times n$ with $0! = 1$)
 $f : \mathcal{X} \mapsto \mathcal{Y}$ a function f with domain \mathcal{X} and codomain \mathcal{Y}
 \lim limit (e.g., unique limit of filter base, function, net, or sequence)
 \rightarrow a limit
 $p_n \rightarrow p$ limit of sequence (p_n)
 $f(x) \rightarrow q$ limit of function f (e.g., as $x \rightarrow p$)
 $f'(x_0+)$ the right-hand derivative of function f at point x_0
 $f'(x_0-)$ the left-hand derivative of function f at point x_0
 $f'(x_0)$ the first (total) derivative of function f at point x_0
 $f''(x_0)$ the second (ordinary) derivative of function f at point x_0
 $f'''(x_0)$ the third (ordinary) derivative of function f at point x_0
 $f^{(n)}(x_0)$ the n^{th} (ordinary) derivative of function f at point x_0 where $n \in \{4, 5, 6, \dots\}$
 $d f / d t$ total derivative of function f at point t
 $d^2 f / d t^2$ second total derivative of function f (i.e., f'')
 $d^3 f / d t^3$ third total derivative of function f (i.e., f''')
 $d^n f / d t^n$ n^{th} total derivative of function f (i.e., $f^{(n)}$)
 $\partial f / \partial x$ partial derivative of function f with respect to x
 $\partial^2 f / \partial x \partial y$ partial derivative of function $\partial f / \partial x$ with respect to y

Vectors and Linear Algebra

y_i	the i^{th} coordinate of vector y
x^\top	the transpose of vector or covector x (i.e., if x is an n -vector then $x = [x_1, x_2, \dots, x_n]^\top$)
A^\top	the transpose of matrix A
e_i	the i^{th} elementary (or standard) basis vector
\mathbb{I}	the identity matrix
$\nabla_x f(x)$	the gradient vector of function f at x
$\nabla_{xx}^2 f(x)$	the Hessian matrix of function f at point x

Probability and Measure Theory

$\mathfrak{B}(\mathcal{U})$	the Borel algebra of set \mathcal{U} (i.e., $\mathfrak{B}(\mathcal{U})$ is the minimal σ -algebra containing the open sets; elements of $\mathfrak{B}(\mathcal{U})$ are called <i>Borel sets</i> and are subsets of \mathcal{U} , so $\mathfrak{B}(\mathcal{U}) \in \mathcal{P}(\mathcal{U})$)
$\int_a^b f(x) dx$	the Lebesgue integral of function f over interval $[a, b] \subset \overline{\mathbb{R}}$ with respect to the Lebesgue measure
$f * g$	convolution of function f with function g (i.e., $(f * g)(t) \triangleq \int_{-\infty}^{\infty} f(\tau)g(t - \tau) d\tau$)
$\delta_a(\mathcal{E})$	Dirac delta measure of set \mathcal{E} at point a (e.g., $f(0) = \int_{-1}^1 f(x)\delta_0(\{x\}) dx$)
$\delta(x - p)$	Simplified Dirac delta measure notation (i.e., $\delta(x - p) \triangleq \delta_p(\{x\})$)
\Pr	Probability measure
$(\mathcal{U}, \Sigma, \Pr)$	Probability space with outcomes \mathcal{U} , σ -field of events Σ , and probability measure \Pr
$\{X \leq a\}$	Measurable set induced by preimage of random variable X (i.e., $\{\zeta \in \mathcal{U} : X(\zeta) \leq a\}$)
$\Pr(X \leq a)$	Probability induced by preimage of random variable X (i.e., $\Pr(\{\zeta \in \mathcal{U} : X(\zeta) \leq a\})$)
$F_X(x)$	Cumulative distribution function for random variable X (i.e., $F_X(a) \triangleq \Pr(X \leq a)$)
$f_X(x)$	Probability density function for random variable X (i.e., $F_X(a) = \int_{-\infty}^a f_X(x) dx$)
$E(X)$	Expectation of random variable X (i.e., $\int_{-\infty}^{\infty} x f_X(x) dx$)
$E(g(X))$	Expectation of function g of random variable X (i.e., $\int_{-\infty}^{\infty} g(x) f_X(x) dx$)
$\text{cov}(X, Y)$	Covariance of random variables X and Y (i.e., $\text{cov}(X, Y) = E(XY) - E(X)E(Y)$)
$F_{XY}(x, y)$	Joint distribution function for random variables X and Y (i.e., $F_{XY}(a, b) \triangleq \Pr(X \leq a, Y \leq b)$)
$f_{XY}(x, y)$	Joint density function for random variables X and Y
$f_{Y X}(y x)$	Conditional density function for random variable Y given $X = x$

$F_{Y X}(y x)$	Conditional distribution function for random variable Y given $X = x$
$E(Y X)$	Conditional expectation of Y given X
$(N(t) : t \in \mathbb{R}_{\geq 0})$	Random process (i.e., $N(t)$ is a random vector for all $t \in \mathbb{R}_{> 0}$)
$Y(t) \xrightarrow{a.s.} Y$	Random process $Y(t)$ converges almost surely (i.e., $\Pr(\lim_{t \rightarrow \infty} Y(t) = Y) = 1$) to Y
$\text{aslim}_{t \rightarrow \infty} Y(t) = Y$	Random process $Y(t)$ converges almost surely (i.e., with probability 1) to Y

INDEX¹

- **A**
- AAV, *see* examples, applications, autonomous vehicle
 - agent, *see* solitary agent model, agent
 - agent model, *see* solitary agent model
 - AGV, *see* examples, applications, autonomous vehicle
 - applications, *see* examples, applications
 - artificial intelligence, [93](#)
 - assumptions, *see* solitary agent model, assumptions
 - automation, *see* examples, applications
 - autonomous air vehicles, *see* examples, applications, autonomous vehicle
 - autonomous ground vehicles, *see* examples, applications, autonomous vehicle
 - average, *see* stochasticity, statistics, expectation
- **B**
- bang-bang control, [47](#)
 - Borel sets, *see* mathematics, sets, Borel algebra
- **C**
- capital budgeting, *see* finance, capital budgets
 - central limit theorem (CLT), *see* stochasticity, central limit theorem
 - central moments, *see* stochasticity, statistics, central moments
 - choice, *see* solitary agent model, choice
 - classical OFT, *see* optimal foraging theory
 - classical optimal foraging theory, *see* optimal foraging theory
 - CLT, *see* stochasticity, central limit theorem
 - cognition, [93](#)
 - combined prey and patch model, *see* combined task-type and processing-length choice problem
 - combined task-type and processing-length choice problem, [11](#), [87](#)
 - compound objectives, *see* optimality, multiobjective optimization
 - Concorde fallacy, *see* sunk-cost effect
 - constraints, *see* optimization constraints
 - contributions
 - biology, [94–96](#)
 - engineering, [93–94](#)
 - costs, *see* solitary agent model, costs
 - currency, *see* solitary agent model, currency
- **D**
- dynamic optimization, [46](#)
- **E**
- ecological rationality, [55](#), [58](#), [59](#)
 - efficiency, *see* excess efficiency
 - EPR, *see* extreme-preference rule

¹Notation: [114](#) is a definition page, and [114](#) is a glossary page.

equilibrium process rate of net gain, *see* equilibrium renewal process rate of net gain

equilibrium renewal process rate of net gain, 42–43

Erlang-2, *see* stochasticity, distributions, Erlang-2

ex post performance measure, 47

examples

- analytical solutions, 90–92
- discounted net gain, 91
- excess efficiency, 92
- rate of excess net gain, 90–91

applications, 93–94

- autonomous vehicle, 1, 4, 31, 94
- military, 94
- queueing, 94
- surveillance, 31
- temperature control, 94

graphical solutions

- excess efficiency, 56–58
- rate of excess net gain, 53–55
- reward-to-variability ratio, 59
- reward-to-variance ratio, 61

excess efficiency, 55–56, 56–58

- graphical optimization, 56–58

expectation, *see* stochasticity, statistics, expectation

expected utility, *see* utility theory

expected value, *see* stochasticity, statistics, expectation

extreme-preference rule (EPR), 84, 106

..... **F**

finance

- capital budgets, 72
- investment returns, 72
- modern portfolio theory (MPT), 33
- portfolios, 72
- post-modern portfolio theory (PMPT), 34, 70–73

- stochastic dominance (SD), 72–73, 106
- first-order (FSD), 72–73, 106
- third-order (TSD), 73, 106

finite expectation, *see* stochasticity, statistics, finite expectation

foraging, *see* optimal foraging theory

foraging model, *see* solitary agent model

frontier, *see* optimality, Pareto efficient

future directions, 70, 96–97

..... **G**

gain, *see* solitary agent model, point gain

gain success threshold, 44–45, 52, 55, 58, 61

gain threshold, *see* gain success threshold

Gaussian distribution, *see* stochasticity, distributions, normal

gross gain success threshold, 55

..... **I**

i.i.d., *see* stochasticity, random variable, i.i.d.

independent and identically distributed (i.i.d.), *see* stochasticity, random variable, i.i.d.

independent random variables, *see* stochasticity, random variable, independent

investment returns, *see* finance, investment returns

..... **K**

Karush-Khun-Tucker (KKT) conditions, *see* optimality, KKT conditions

KKT conditions, *see* optimality, KKT conditions

Kuhn-Tucker conditions, *see* optimality, KKT conditions

..... **L**

Lagrange multiplier method, *see* optimality, Lagrange multiplier method

location, *see* stochasticity, distributions, location-scale family

location-scale, *see* stochasticity, distributions, location-scale family

long-term average rate of net gain, **19, 27, 39–44**

- graphical optimization, **43–44**
- justification, **39–40**
- limit, **40–41**
- opportunity cost, **41–42**

long-term rate of net gain, *see* long-term average rate of net gain

lower-partial moment, *see* stochasticity, statistics, lower-partial moment

lower-partial variance, *see* stochasticity, statistics, lower-partial variance

LPM, *see* stochasticity, statistics, lower-partial moment

LPV, *see* stochasticity, statistics, lower-partial variance

..... **M**

marginal value theorem (MVT), **44, 87, 106**

Markov renewal cycle

- OFT cycle, **108**

Markov renewal cycles

- OFT cycle, *see* solitary agent model, classical analysis, OFT cycle
- processing cycle, *see* solitary agent model, processing-only analysis, processing cycle

Markov renewal process, *see* stochasticity, Markov renewal process

mathematics, **109–113**

- definition (\triangleq), **109**
- equality ($=$), **109**
- functions, **111**
 - convolution ($*$), **99, 112**
 - gradient (∇), **112**
 - Hessian (∇^2), **112**
 - integral (\int), **71, 112**
 - partial derivative (∂), **111**
 - total derivative (d), **111**

limit (lim or \rightarrow), **111**

logic, **111**

- equivalence (\iff), **111**
- implication (\implies), **111**

n -tuple, **110**

numbers, **109–110**

- Euclidean n -space (\mathbb{R}^n), **109**
- extended real numbers ($\overline{\mathbb{R}}$), **109**
- integers (\mathbb{Z}), **109**
- natural numbers (\mathbb{N}), **109**
- rationals (\mathbb{Q}), **109**
- real negative numbers ($\mathbb{R}_{<0}$), **109**
- real non-negative numbers ($\mathbb{R}_{\geq 0}$), **109**
- real non-positive numbers ($\mathbb{R}_{\leq 0}$), **109**
- real non-zero numbers ($\mathbb{R}_{\neq 0}$), **109**
- real number intervals, **110–111**
- real numbers (\mathbb{R}), **109**
- real positive numbers ($\mathbb{R}_{>0}$), **109**
- whole numbers (\mathbb{W}), **109**

order, **111**

- infimum (inf), **111**
- maximum (max), **111**
- minimum (min), **111**
- supremum (sup), **111**

ordered pair, **110**

sets, **110–111**

- Borel algebra (\mathfrak{B}), **112**
- Cartesian product (\times), **110**
- empty set (\emptyset), **110**
- exclusion (\notin), **110**
- power set (\mathcal{P}), **110**
- set complement (c), **110**
- set difference ($-$), **110**
- set element (\in), **110**
- set intersection (\cap), **110**
- set union (\cup), **110**
- subset (\subseteq or \subset), **110**

- superset (\supseteq or \supset), 110
- vector spaces, 111–112
 - matrix transpose (T), 112
 - vector transpose (T), 112
- maximal expected utility, *see* utility theory
- maximal utility, *see* utility theory
- mean, *see* stochasticity, statistics, expectation
- mean-lower-partial-moment (MLPM) analysis, 71
- mean-lower-partial-variance analysis, 71
- mean-semivariance analysis, *see* mean-lower-partial-moment analysis, *see* mean-lower-partial-variance analysis
- mean-variance analysis (MVA), 44–47, 70, 106
- merge before split approach, 13
- merged Poisson process, *see* stochasticity, Poisson process, merged
- MLPM, *see* mean-lower-partial-moment analysis
- model, *see* solitary agent model
- model assumptions, *see* solitary agent model, assumptions
- model weaknesses, *see* solitary agent model, weaknesses
- modern portfolio theory (MPT), *see* finance, modern portfolio theory
- moments, *see* stochasticity, statistics, moments
- MPT, *see* finance, modern portfolio theory
- MSA, *see* mean-lower-partial-moment analysis
- multiobjective optimization, *see* optimality, multiobjective optimization
- mutually independent and identically distributed (i.i.d.), *see* stochasticity, random variable, i.i.d.
- MVA, *see* mean-variance analysis
- MVT, *see* marginal value theorem
- **N**
- net gain success threshold, *see* gain success threshold
- net point gain, *see* solitary agent model, net point gain
- normal distribution, *see* stochasticity, distributions, normal
- numbers, *see* mathematics, numbers
- **O**
- OFT, *see* optimal foraging theory
- OFT cycle, *see* solitary agent model, classical analysis, OFT cycle
- opportunity cost, *see* long-term average rate of net gain, opportunity cost
- optimal foraging theory, 1, 3, 5, 12–22, 27, 106
- optimality
 - compound objectives, *see* optimality, multiobjective optimization
 - KKT conditions, 36
 - Lagrange multiplier method, 36, 38
 - multiobjective optimization, 35–36
 - linear combination, 36
 - maximin, 36
 - Pareto efficient, 36, 36
 - Pareto frontier, *see* optimality, Pareto efficient
 - Pareto optimal, *see* optimality, Pareto efficient
- optimization constraints, 36–38, 66–69
 - encounter-rate, 37
 - gain and cost, 68
 - gain and time, 67
 - mean and standard deviation, 69
 - mean and variance, 68–69
 - nutrients, 37
 - time, 36
- **P**

- Pareto efficient, *see* optimality, Pareto efficient
 - Pareto frontier, *see* optimality, Pareto efficient
 - Pareto optimal, *see* optimality, Pareto efficient
 - Pareto tradeoffs, 63–66
 - efficiency, 64
 - gain discounted by cost, 64
 - gain and time, 64
 - gain discounted by time, 64
 - mean and standard deviation, 65
 - mean discounted by standard deviation, 65
 - mean and variance, 65–66
 - mean discounted by variance, 65
 - patch model, *see* task processing-length choice problem
 - patch overstaying, 48
 - rational explanation, 54
 - PMPT, *see* finance, post-modern portfolio theory
 - point gain, *see* solitary agent model, point gain
 - Poisson process, *see* stochasticity, Poisson process
 - portfolios, *see* finance, portfolios
 - post-modern portfolio theory (PMPT), *see* finance, post-modern portfolio theory
 - prey model, *see* task-type choice problem
 - probability, *see* stochasticity
 - probability measure, *see* stochasticity, probability measure
 - probability of success, 45
 - probability space, *see* stochasticity, probability space
 - processing, *see* solitary agent model, processing cycle, *see* solitary agent model, processing-only analysis, processing cycle
 - pseudo-deterministic, *see* stochasticity, random variable, pseudo-deterministic
- **Q**
- queueing, *see* examples, applications, queueing
- **R**
- random variables, *see* stochasticity, random variables
 - randomness, *see* stochasticity
 - rate maximization, 39–44, 52
 - rate of excess net gain, 52, 52–55
 - analytical optimization, 90–91
 - graphical optimization, 52–55
 - rational agent, *see* utility theory
 - rational objective function, 74–90
 - optimal solution
 - constant disadvantage, 85–87
 - decreasing advantage-to-disadvantage, 87–90
 - rationality, *see* ecological rationality
 - real numbers, *see* mathematics, numbers, real numbers
 - renewal process, *see* stochasticity, Markov renewal process
 - reward-to-variability
 - graphical optimization, 47
 - reward-to-variability ratio, 46, 47, 58, 58, 58–59
 - graphical optimization, 58–59
 - reward-to-variance ratio, 61, 61, 61
 - graphical optimization, 61
 - risk minimization, *see* risk sensitivity
 - risk sensitivity, 44–47, 52, 65–66
 - robotics, *see* examples, applications
- **S**

satisficing, **49**
 scale, *see* stochasticity, distributions,
 location-scale family
 SD, *see* finance, stochastic dominance
 searching, *see* solitary agent model,
 searching
 semivariance, *see* stochasticity, statistics,
 lower-partial variance
 Sharpe ratio, *see* reward-to-variability ratio
 skewness, *see* stochasticity, statistics,
 skewness
 skewness preference, **35**
 solitary agent model, **3–5**, **6–12**, **12–32**
 agent, **4**
 assumptions, **6–8**
 choice, **4**
 classical analysis, **13–22**
 limits, **19**
 OFT cycle, **16**, **16**
 optimization, **39–49**, **74–92**
 Poisson encounters, **13–14**
 renewal process, **16–17**
 reward processes, **14–15**, **18**
 statistics, **14–16**, **18–22**
 variance, **20–22**
 comparison of analyses, **29–30**
 costs, **4**
 currency, **4**
 decision variables, **11**
 net point gain, **4**
 parameters, **9–11**, **106–108**
 point gain, **4**
 processing, **4**
 processing-only analysis, **5**, **22–28**
 limits, **26–27**
 optimization, **49–69**, **74–92**
 Poisson encounters, **22–23**
 processing cycle, **25**, **25**
 renewal process, **25**
 reward processes, **23–26**
 statistics, **24**, **26–28**
 variance, **27–28**
 random processes, **8–9**, **12**
 searching, **4**
 tasks, **4**
 weaknesses, **30–32**
 speed-accuracy tradeoff, **38**
 split before merge approach, **22**
 split Poisson process, *see* stochasticity,
 Poisson process, split
 stochastic dominance, *see* finance,
 stochastic dominance
 stochastic model, *see* solitary agent model
 stochasticity, **3**, **112–113**
 central limit theorem (CLT), **46**, **106**
 distributions, **72**
 double exponential, **45**
 Erlang-2, **99**, **99**
 exponential, **45**
 location, *see* stochasticity, distribu-
 tions, location-scale
 location-scale family, **45**, **45–46**
 log-normal, **45**
 normal, **45**, **46**, **46**
 scale, *see* stochasticity, distribu-
 tions, location-scale
 symmetry, **34**, **46**
 uniform, **45**
 Markov renewal process, **16–17**, **98–99**
 limits, **98–99**
 Markov renewal-reward process, **18**,
 26
 Poisson process, **98**, **99**
 merged, **14**, **15**, **22**, **23**
 split, **22**, **23**
 probability measure, **6**, **112**
 probability space, **6**, **98**, **112**
 random process, **113**
 almost sure limit (aslim or $\xrightarrow{a.s.}$),
 113
 random variable, **6**, **34**, **72**

cumulative distribution function (F), 112

i.i.d., 8, 106

independence, 7, 12

probability density function (f), 112

pseudo-deterministic, 7–8, 12, 20, 27

uncorrelated, 7

statistics

- central moments, 20, 34
- conditional expectation, 113
- expectation (E), 6, 112
- expectation of function, 112
- expected value, *see* stochasticity, statistics, expectation
- finite expectation, 8, 8, 12
- lower-partial moment (LPM), 71, 106
- lower-partial variance (LPV), 70–71, 106
- mean, *see* stochasticity, statistics, expectation
- moments, 6
- normal, 46
- semivariance, *see* stochasticity, statistics, lower-partial variance
- skewness, 34–35, 45
- standard deviation, 44, 44
- variance (var), 20

success threshold, *see* gain success threshold

sunk-cost effect, 55, 58

survival threshold, *see* gain success threshold

symmetry, *see* stochasticity, distributions, symmetry

..... **T**

task processing-length choice problem, 11

task-type choice problem, 11, 86

tasks, *see* solitary agent model, tasks

threshold, *see* gain success threshold

..... **U**

UAV, *see* examples, applications, autonomous vehicle

UGV, *see* examples, applications, autonomous vehicle

uncorrelated random variables, *see* stochasticity, random variable, uncorrelated

unmanned air vehicles, *see* examples, applications, autonomous vehicle

unmanned ground vehicles, *see* examples, applications, autonomous vehicle

utilitarianism, *see* utility theory

utility function, *see* utility theory

utility theory, 65–66, 72–73

..... **V**

variance, *see* stochasticity, statistics, variance

vehicles, *see* examples, applications, autonomous vehicle

von Neumann-Morgenstern utility, *see* utility theory

..... **Z**

z -score, *see* reward-to-variability ratio

zero-one rule, 84

PEOPLE¹

- Andrews, Burton W., **1**, 94, **100**, *103*
Arkes, Hal R., **55**, **100**
Ayton, Peter, **55**, *100*
- Bateson, Melissa, **49**, **50**, **100**
Bawa, Vijay S., **45**, **70–73**, **100**
Ben-Haim, Yakov, **49**, *101*
Bernhard, Pierre, **51**, *105*
Bernoulli, Jacob, **9**, **10**
Bertsekas, Dimitri P., **36**, **67**, **77**, **100**
Bodurtha Jr., James N., **vi**
Boivin, Guy, **51**, *105*
- Caraco, Thomas, **66**, **73**, **101**
Carmel, Yohay, **49**, **101**
Chamberlain, Gary, **45**, **101**
Charnov, Eric L., **3**, **5**, **13**, 38–42, **44**,
45, 47, 49, **50**, **87**, **101**, *103*, *104*
Chunhachindaa, Pornchai, **34**, **101**
Clark, Colin W., **51**, **101**, *102*
Cody, Martin L., **31**, **101**
Cruz Jr., Jose B., **vi**
- Dandapanib, Krishnan, **34**, *101*
Darwin, Charles R., **39**
Descartes, René, **110**
Downhower, Jerry F., **vi**
- Ebie, Jessica D., **v**
Émile Borel, **6**, **112**
Engen, Steinar, **41**, **101**
Euclid, **109**
Euler, Leonhard, **109**
- Falkner, Neil, **vi**
Ferguson, Kathleen W., **70**, **71**, *104*
- Gauß, Carl Friedrich, **46**
Gendron, Robert P., **37**, **101**
Gigerenzer, Gerd, **49**, *102*
Grootveld, Henk, **71**, **101**
- Hall, Donald J., **3**, *105*
Hallerbach, Winfried, **71**, *101*
Hamelin, Frédéric, **51**, *105*
Hamidb, Shahid, **34**, *101*
Harder, Lawrence D., **49**, **50**, **101**
Hesse, Ludwig Otto, **78**, **81–83**, **112**
Higashi, Masahiko, **49**, *102*
Houston, Alasdair I., **31**, **41**, **42**, **51**, **64**,
101
Hughes, Roger N., **3**, **102**
Hutchinson, John M. C., **49**, **102**
- Iwasa, Yoh, **49**, **102**
- Johns, M. V., **19**, **99**, **102**
- Kacelnik, Alex, **49**, **50**, *100*
Kane, Alex, **34**, **102**
Karush, William, **36**
Konno, Hiroshi, **34**, **102**
Krebs, John R., **1**, **5**, 11–13, **41**, **43**, **46**,
47, **50**, **66**, **73**, **84**, **86–88**, **91**, **92**,
104
Kuhn, Harold W., **36**

¹Notation: **121** is a first-author citation, and *121* is a bibliography page.

- Lagrange, Joseph-Louis, **36, 38, 67, 77, 78, 81, 83**
- Lai, Tsong-Yue, **34, 102**
- Lawlor, Lawrence R., *49, 104*
- Lebesgue, Henri L., **6, 34, 46, 71, 98, 112**
- Lee, Wayne Y., **71, 102**
- Lindenberg, Eric B., *71, 100*
- Lintner, John, **70, 102**
- Mangel, Marc, **51, 101, 102**
- Mao, James C. T., **71, 102**
- Markov, Andrei A., **14, 16–18, 24–26, 98, 107, 108**
- Markowitz, Harry M., **1, 61, 65, 70, 102**
- Marschak, Jacob, **66, 103**
- McNamara, John M., *31, 41, 42, 47, 51, 64, 101, 103*
- Meyer, Jack, **45, 103**
- Miller, R. G., Jr., *19, 99, 102*
- Morgenstern, Oskar, *66, 72, 104*
- Moses, Randolph L., **vi**
- Mossin, Jan, **70, 103**
- Murdoch, William W., **37, 103**
- Nobel, Alfred B., **46**
- Nonacs, Peter, **48, 50, 103**
- Orians, Gordon H., *3, 5, 13, 42, 44, 50, 87, 101*
- Owen, Joel, **45, 103**
- Papoulis, Athanasios, **99, 103**
- Pareto, Vilfredo, **36, 41, 42, 50, 53, 54, 56–69**
- Passino, Kevin M., **v, 1, 10, 37, 94, 100, 103**
- Pavlic, Eileen M., **v**
- Pavlic, Paul V., **v**
- Pavlic, Theodore P., **1, 10, 37, 103**
- Pillai, S. Unnikrishna, *99, 103*
- Poisson, Siméon-Denis, **8, 10, 13–15, 22, 23, 30, 94, 98, 99, 107, 108**
- Prakash, Arun J., *34, 101*
- Pulliam, H. Ronald, **3, 37–40, 50, 103**
- Pyke, Graham H., **38–40, 50, 103**
- Quijano, Nicanor, **1, 94, 103**
- Rabinovitch, Ramon, *45, 103*
- Rao, Ramesh K. S., *71, 102*
- Rapport, David J., **36–39, 53, 103**
- Real, Leslie A., **49, 50, 65, 66, 73, 101, 104**
- Ridley, Matt, **vi**
- Rom, Brian M., **70, 71, 104**
- Rothschild, Michael, **45, 104**
- Schoener, Thomas W., **3, 31, 38–40, 50, 64, 104**
- Sharpe, William F., **46, 47, 58, 61, 70, 104**
- Shirakawa, Hiroshi, *34, 102*
- Staddon, John E. R., *37, 101*
- Stenseth, Nils Chr., *41, 101*
- Stephens, David W., **1, 5, 11–13, 39, 41, 43, 45–47, 49, 50, 66, 73, 84, 86–88, 91, 92, 104**
- Stiglitz, Joseph E., *45, 104*
- Templeton, Alan R., **49, 104**
- Timmons, Kenneth, **iv**
- Tobin, James, **61, 65, 70, 104**
- Tsiang, S. C., **34, 104**
- Tucker, Albert W., *36*
- Viniotis, Yannis, **14, 104**
- von Neumann, John, **66, 72, 104**
- Waite, Thomas A., **v, 1, 100**
- Wajnberg, Eric, **51, 105**
- Werner, Earl E., **3, 105**
- Whitehead, Siân C., *49, 100*
- Yamamura, Norio, *49, 102*
- Yamazaki, Hiroaki, *34, 102*
- Zheng, Yuan F., **vi**