Three Essays on Industrial Organization

Dissertation

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

By
Jafar Olimov

Graduate Program in Agricultural, Environmental, and Development Economics
The Ohio State University
2013

Dissertation Committee:
Professor Brian Roe, Advisor
Professor John Kagel
Professor James Peck
Professor Ian Sheldon

Abstract

In this dissertation I study three topics in Industrial Organization. In the first chapter I characterize a set of subgame perfect Nash equilibria in the alternating-offers bargaining game with one-sided outside option and differentiated depreciation rates. I use this set of equilibria to solve the hold-up problem. In particular, I argue that specific investment imposes a cost on the investor's partner, because specific investment lowers the likelihood that the investor's partner extracts an additional payoff above her outside option in the set of equilibria which we characterize. If for the investor's partner the expected marginal cost of specific investment is equal to the expected marginal benefit, then there is no ex ante under-investment and the hold-up problem is completely resolved.

In the second chapter I study seller behavior using data from eBay auctions of used tractors. I relax the standard assumption that sellers know the distribution functions of items' valuations and find that uninformed and patient sellers use secret reserve prices to run unsuccessful eBay auctions to learn parameters of these unknown distribution functions. I find that secret reserve prices have strong positive effect on sale prices. I provide a novel theoretical justification for the use of secret reserve prices and show that eBay serves not only as a selling platform but also as an affordable value-appraising mechanism for items whose valuation is not easily available or is costly to obtain.

In the third chapter I study a model of quality sorting between electronic and physical platforms. In the model a seller can auction an item with both opaque and transparent quality attributes in either platform. Bidders can observe perfectly the quality of both the transparent and the opaque attribute in a physical platform. In an electronic platform bidders can observe perfectly only the quality of the transparent attribute but not the opaque attribute. I use Spence's signaling model (1973) to derive seller's equilibrium listing strategy. I find that conditional on the quality of the transparent attribute, the quality of the opaque attribute in a physical platform is always no worse than the quality of the opaque attribute in an lectronic platform. I also find that when items on sale have both transparent and opaque attributes, it is impossible to compare the item's overall quality in both platforms without restricting bidders' beliefs about the quality of the opaque attribute in an electronic platform or without introducing an additional structure into the model. The main conclusion is that it is not always true that the overall quality of items in an electronic platform is necessarily lower than the overall quality of items in a physical platform.

Acknowledgements

I would like to thank my adviser Brian Roe for his assistance, time, and expertise in preparation of this manuscript. I would also like to thank my wife Natalia, without whose help this dissertation would not be possible.

Vita

January 1, 1982	Born, Moscow, USSR
2004	BA Economics, Denison University
2008	MA Economics, The Ohio State University

Publications

B. Roe, T. Wyszynski and J.M. Olimov, "Pigs in Cyberspace: A Natural Experiment Testing Differences between Online and Offline Club-Pig Auctions," *American Journal of Agricultural Economics*, (2011) 93 (5): 1278-1291;

Fields of Study

Major Field: Agricultural, Environmental and Development Economics

Table of Contents

List of Figures

List of Tables

Chapter 1

Optimal Specific Investment

Incomplete contract theory admits that contracts either cannot specify all contingencies or that some observable actions are not verifiable by a third party. Consequently, after a contract is signed and uncertainty is realized, contracting parties may renegotiate the contract conditions given their ex post bargaining power. An investor signing an incomplete contract anticipates this ex post renegotiation and makes an individually optimal choice of investment ex ante. In other words, the investor ignores the benefit of her partner(s) and under-invests from the socially optimal point of view. This problem is commonly known as the hold-up problem[1]. The problem of under-investment or the hold-up problem is particularly strong in incomplete contracts with relationship-specific (hereafter, specific) investment. As opposed to relationships with generic investment, in a relationship with specific investment the value of investment is partially or completely lost if the relationship is terminated. As a result, an investor has particularly strong incentives to ignore the benefit of her partners, given that she cannot recoup her investment if the relationship is terminated ex post.

In this chapter we propose a solution to the hold-up problem. We argue that specific investment increases the likelihood of a lower negotiation payoff obtained by the investor's partner. If for the investor's partner the expected cost associated with the increase in the

[1] The hold-up problem is similar in nature to the problem of under-provision of a public good. In the context of the hold-up problem, all parties benefit from a specific investment, but only the investor bears the associated costs. The investor maximizes his own benefit without regard for his partners and under-invests from the social point of view. The investor is reluctuant to take into account the benefit of his partners, because he is afraid that his partners may renegotiate trade terms after the investor makes the individually supra-optimal investment.

likelihood of a lower negotiation payoff is equal to the expected benefit from specific investment, then the investor's choice of specific investment is first-best optimal, and the hold-up problem is completely resolved.

The resolution of the hold-up problem in our model depends on two key elements: the presence of an additional payoff above the outside option of the investor's partner, which the investor's partner can lose as a result of increase in specific investment, and the effect of specific investment on the distribution of negotiation outcomes. In our model the investor's partner can obtain a payoff above her outside option, because her outside option (which represents an alternative tradable good produced with generic investment) has a lower depreciation rate than the tradable good with specific investment in the relationship. Since specific investment raises the value of the tradable good in the relationship but has no effect on the distribution of the generically produced outside option, with an increase in specific investment the investor's partner is less likely to use her outside option as a credible threat in the negotiation process and is less likely to extract an additional payoff above her outside option. The size and the presence of this additional payoff depends on the depreciation rates on the tradable good in the relationship and the outside option.

As opposed to other approaches to the hold-up problem, we do not rely on any contractual or third-party enforcement mechanisms. Moreover, the resolution of the hold-up problem in our approach is not driven by the investor's incentive to invest but rather by the expected cost of specific investment imposed on the investor's partner. Depending on parameter values it is possible that the investor's choice of investment results in over-investment or under-investment, if the expected marginal cost imposed on the investor's partner either exceeds or falls below the expected marginal benefit from specific investment.

Coase (1937) is one of the first to discuss the hold-up problem in the context of firm boundaries. In particular, the hold-up problem has been considered a key justification for vertical integration. Grout (1984) offers a formal treatment of the hold-up problem in the context of incomplete contracts and finds that the investor necessarily under-invests. Grossman and Hart (1986) and Hart and Moore (1988) offer another formalization of the hold-up problem and arrive at a similar conclusion: if specific investment is super-modular and contracts are incomplete, then the investor under-invests. Grossman and Hart (1986) also suggest a solution to the hold-up problem, in which a party undertaking specific investment owns the means of production. In their setup the hold-up problem is resolved because the investing party gains residual control rights, appropriates the full surplus from his actions, and necessarily undertakes the first-best level of specific investment. Sloof, Sonnemans, and Oosterbeek (2004) propose a similar solution, where the investor's partner is exogenously provided a binding non-random outside option. As a result, the investor becomes a residual claimant with sufficient incentives to invest optimally.

A number of other solutions to the hold-up problem have been suggested as well, however all of them employ some form of contractual solutions, third-party enforcement of contract terms, or breach remedies. For example, Aghion, Dewatripont, and Rey (1994) propose a solution where trade partners are offered a renegotiable contract specifying a "financial hostage" for the buyer and a breach remedy for the seller. Nöldeke and Schmidt (1995) employ an option contract involving a third-party verifiability of the delivery decision, while Edlin and Reichelstein (1996) introduce a contract specifying breach remedies or expectation damages contingent on specific performance.

Several recent studies employ dynamic solutions to the hold-up problem. Che and Sákovics (2004) find that the allocation of bargaining rights loses its importance in a dynamic setting. If agents are willing to participate in a relationship and they are sufficiently patient, the investment dynamics alone are sufficient to induce optimal specific investment. However, this result holds only in the infinite-time setting. In the finite-time setting, the investing party under-invests. Guriev and Kvasov (2005) propose a solution to the hold-up problem in the finite-time setting. The authors employ a modification of the Nash folk theorem for finitely repeated games and find that a fixed-term contract of a sufficient length can induce the investor to undertake the first-best optimal specific investment.

The major drawback of the existing approaches to the hold-up problem is that they rely on contractual solutions, third parties, or continuous investing. Practically, such solutions are at odds with the original assumption that the complexity of the economic environment is what gives rise to incomplete contracts. In this respect our approach is different, because we do not rely on any complex mechanisms. We employ a simple complete information bilateral trading model with static one-sided investment and no verifiability. The parties in our model split the trade surplus according to Rubinstein's alternating-offers bargaining game with one-sided outside option. As a result, in our solution there is no internal inconsistency between assumptions underlying the problem and the solution to the problem.

In the next section we discuss empirical evidence on the hold-up problem. In section 1.2 we formalize our model and characterize all equilibria in the bargaining game used in the chapter. In section 1.3 we characterize the first-best level of specific investment. In section 1.4 we define the equilibrium level of specific investment. In section 1.5 we introduce conditions under which the equilibrium specific investment is equal to the first-best specific investment and provide a specific example of the investment game with uniformly distributed outside option. In section 1.6 we discuss results and present testable hypotheses arising from our model. In the last section we conclude.

1.1 Empirical Evidence

The major conclusion from the literature is that the hold-up problem may be one of the main reasons for vertical integration (see, for example, Williamson (1985)). A number of empirical works find that the presence of specific investment indeed increases the likelihood of vertical integration (see Lafontaine and Slade (2007) for an overview). However, there are many examples of contractual relationships with specific assets and no vertical integration. The most closely studied example of a relationship involving specific investment and independent ownership is the Japanese auto industry. As opposed to North American car manufacturers, Japanese auto-makers rely on supplies produced by independently owned manufacturers (see Aoki (1990) for more details). Although there is a high potential for Japanese inputs suppliers to under-invest out of fear of hold-up, the duration and stability of relationships between Japanese manufactures and their inputs suppliers suggests optimality of investment.

Another example of a non-integrated relationship with specific investment is airline alliances. For example, Star Alliance currently includes 27 major airline companies which share each other's networks and provide ticketing and other services to customers from other members of the alliance. We can expect that because of the hold-up problem members of Star Alliance should under-invest in the jointly used facilities and services, which eventually should lead to the dissolution of the alliance or to mergers under a single ownership. Nevertheless, we observe neither the dissolution nor mergers among members of the alliance. Moreover, the number of participating non-integrated members in the alliance has been increasing ever since its inception in 1997.

Holmström and Roberts (1998) report that electronics and software industries can serve as another case of relationships with specific investment where joint ownership is relatively rare. For example, the current relationship between Foxconn International Holdings Ltd., a Taiwanese manufacturer of iPads and iPhones, and Apple Inc. is illustrative of a relationship with highly specific investment and independent ownership. In fact, Foxconn produces parts for almost all major high-tech firms such as Microsoft, Dell, Cisco, Nokia, Sony-Ericsson etc.

To put our discussion in a well-known context, let's consider the famous GM-Fisher Body example of the hold-up problem. Klein, Crawford and Alchian (1978) and Klein (2000) report that in 1919 GM signed a 10-year contract with Fisher Body, then the largest producer of wood-based closed car bodies[2]. Since Fisher Body had to make extensive investment in car body dies and presses specifically for GM's cars and these dies and presses were nearly worthless outside the relationship, the GM-Fisher Body relationship involved specific investment. The pricing of car bodies was based on the cost-plus formula with GM completely covering Fisher's production costs in addition to paying a 17.6% upcharge payment. The GM-Fisher relationship worked well from 1919 to 1924; however, because the demand for cars with closed bodies unexpectedly rose in the early 1920's, GM demanded to revise the contract and increase the production of closed car bodies by Fisher Body. The Fisher brothers, the owners of Fisher Body, raised production but were able to retain the old pricing scheme and obtain "greater than competitive costs and prices...to earn a greater than competitive return on capital."[3] Thus, the Fisher brothers were able to get a higher share of the total surplus by keeping the 17.6% upcharge payment in addition to lowering the production costs through economies of scale. Klein (2000) claims that Fisher Body refused to construct a new body plant close to GM's production facility. As a result, in 1925 GM acquired the remaining 40% of Fisher Body's shares and vertically integrated the producer of car bodies.

The opinions of researchers diverge on the reason why GM acquired Fisher Body. While Klein (2000) claims that the refusal to locate Fisher's plant close to GM's manufacturing facility was an instance of a hold-up, and GM acquired the Fisher Body to remove such issues in the future, Coase (2000), Freeland (2000), and Casadesus-Masanell and Spulber (2000) agree that the acquisition of the Fisher Body by GM took place to retain the Fisher brothers in GM and prevent their departure to GM's competitors such as Ford.

Although the actual reason for integration between GM and Fisher Body is not so important, other details of the account of the GM-Fisher Body relationship can highlight

[2]According to Casadesus-Masanell and Spulber (2000) and Freeland (2000), Fisher Body was the major producer of wood-based composite closed car bodies and supplied closed car bodies to all major car producers including GM and Ford. Other closed car body producers relied on a different technology which involved all-metal car body production.

[3]Klein (2000), pg. 115.

the mechanism proposed in this chapter. The example of the GM-Fisher Body relationship shows that by undertaking investment in GM-specific production facilities, the Fisher brothers lowered the anticipated ability of GM to credibly use the termination threat in the negotiation process, as it would be harder for GM to find an alternative trade partner with a similar production technology in case of a departure from Fisher Body. In addition, after undertaking specific investment, the Fisher brothers were able to negotiate a larger share of the trade surplus with GM: the Fisher brothers managed to retain the cost-plus pricing formula intact after the original contract was renegotiated. Hence, the GM-Fisher Body case shows that GM in fact experienced a loss as a result of Fisher Body's specific investment.

Moreover, it is unlikely that Fisher Body under-invested in GM-specific production facilities given that exactly this sort of investment weakened GM's bargaining power and allowed the Fisher brothers to negotiate better trade terms with GM. According to Casadesus-Masanell and Spulber (2000), over the years when Fisher Body cooperated with GM as an independent contractor, Fisher Body's annual output rose more than 3 fold from 134,767 closed car bodies in 1919 to 574,979 in 1924. Thus, factually there is no evidence of under-investment in the GM-Fisher case, and the account of the GM-Fisher Body relationship tells that Fisher Body probably over-invested in GM-specific facilities. For example, White reports that GM was the last US car manufacturer to switch from the wood-based composite body production (the specialization of Fisher Body) to a more technologically advanced all-metal car body production, and that GM continued to rely on the wood-based technology until 1937[4].

The account of the GM-Fisher Body relationship shows little evidence of the hold-up problem, and GM's decision to acquire Fisher Body could have been induced by reasons other than the need to eliminate under-investment. In the rest of the chapter we attempt

[4] In Casadesus-Masanell and Spulber (2000), p. 85.

to explain the salient features of the GM-Fisher Body story such as the ability of the Fisher brothers to negotiate better trade terms and show how this ability could have resolved the hold-up problem.

1.2 Model

In this section we study a static bilateral trading model of specific investment. Two risk-neutral contracting parties (a buyer and a seller) engage in trade over a good whose production involves specific investment. At date $t = 0$ the seller makes specific investment σ, $\sigma \in [0, \Sigma]$, at cost $c(\sigma)$, where $c'(\sigma) > 0$. The seller's specific investment σ lowers the cost of production of the tradable good $s(\sigma)$, $s'(\sigma) < 0$, and raises the buyer's valuation of the tradable good $b(\sigma)$, $b'(\sigma) > 0$.

In the model, the uncertainty is realized in the form of the arrival of the outside option. Since investment σ is relationship-specific and the seller's tradable good has no value outside the relationship, the seller's outside option is always equal to zero. The buyer's outside option is a random variable $\underline{V}$ with cdf $F_{\underline{V}}(.)$, support $[0, V]$, and a realized value $\underline{v}$. The buyer's outside option represents the buyer's valuation of an alternative tradable good produced with a generic investment.

The buyer's outside option arrives only once after the seller makes specific investment and stays available to the buyer during the course of the whole relationship. Since investment σ is relationship-specific, investment σ does not impact the distribution function of the buyer's outside option $F_{\underline{V}}(.)$.

After the uncertainty is resolved and the outside option has arrived, the seller and the buyer negotiate over the division of the trade surplus $w(\sigma) = b(\sigma) - s(\sigma)$. The negotiation process is costly to the buyer and to the seller. With each round of negotiation the trade surplus in the relationship $w(\sigma)$ depreciates at rate δ_1 and the buyer's realized outside option $\underline{v}$ depreciates at rate δ_2. If the seller and the buyer agree on the division of the trade surplus in the negotiation stage, the seller produces the tradable good of quantity $q = 1$ at cost $s(\sigma)$ in addition to facing the cost of specific investment $c(\sigma)$, sells the good to the buyer, and splits the trade surplus with the buyer according to the negotiated division. If the negotiation process breaks down, then the seller faces the cost of specific investment $c(\sigma)$ and does not produce the tradable good ($q = 0$), the buyer takes the realized outside option $\underline{v}$ without any penalty, while the seller does not receive any compensation or breach remedy from the buyer or from any outside party. Hence, whether the trade occurs or not, the seller always covers the cost of specific investment $c(\sigma)$. The hold-up problem occurs, because the seller under-invests ex ante anticipating the loss in her share of the trade surplus due to arrival of a competitive outside option to the buyer. The timing of the model is presented in Figure 1.1.

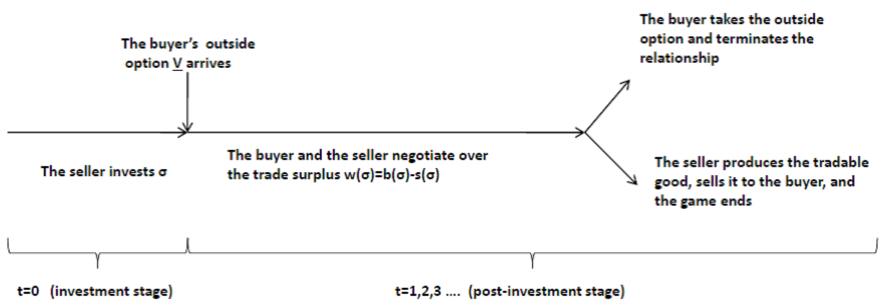

Figure 1.1. Timing of the Model

Since specific investment σ directly raises the buyer's valuation of the tradable good $b(\sigma)$ in addition to lowering the seller's cost of production of the tradable good $s(\sigma)$, in the terminology of Che and Hausch (1999) specific investment σ is a "cooperative investment."

Since our model does not rely on any contractual elements or the verifiability of any parameters of interest, we assume that the realized outside option $\underline{v}$, the quantity of the tradable good q, the seller's cost of the tradeable good $s(\sigma)$, the buyer's valuation of the tradeable good $b(\sigma)$, the level of specific investment σ, and the cost of specific investment $c(\sigma)$ are all observable but not verifiable.

1.2.1 Payoffs

The seller's and the buyer's payoffs should satisfy the parties' participation constraints. In particular, to engage in the production of the tradable good, the payment to the seller's should at least cover the seller's cost of production of the tradable good. Similarly, for the buyer to participate in the relationship, the payment from the buyer should not exceed the buyer's valuation of the tradable good.

Let a be the seller's share of the trade surplus $w(\sigma)$. Then the total payment to the seller is $s(\sigma) + a$, and this payment should satisfy the seller's participation constraint of $s(\sigma) + a \geq s(\sigma)$ or $a \geq 0$. Given the payment to the seller of $s(\sigma) + a$, the buyer's participation constraint requires that $s(\sigma) + a \leq b(\sigma)$ or $a \leq w(\sigma)$. With these constraints in place, we conclude that the seller's share of the trade surplus is $a \in [0, w(\sigma)]$.

If $a = w(\sigma)$, the seller appropriates the whole trade surplus and the buyer obtains none of it. Consequently, the seller's payoff in the relationship is $U^S = -c(\sigma) + [(s(\sigma) + a) - s(\sigma)]q = -c(\sigma) + w(\sigma),$ while the buyer's payoff is $U^B = [b(\sigma) - (s(\sigma) + a)]q = 0$.

If $a = 0$, then the buyer appropriates the whole trade surplus while the seller gets none of it. As a result, the seller's payoff from the relationship is $U^S = -c(\sigma) + [(s(\sigma) + a) - s(\sigma)]q = -c(\sigma)$ and the buyer's payoff from the relationship is $U^B = [b(\sigma) - (s(\sigma) + a)]q = w(\sigma)$.

If $a \in (0, w(\sigma))$, then the seller's payoff is $U^S = -c(\sigma) + [(s(\sigma) + a) - s(\sigma)]q = -c(\sigma) + a$ and the buyer's payoff is $U^B = [b(\sigma) - (s(\sigma) + a)]q = w(\sigma) - a$.

The joint payoff of both the seller and the buyer is $U^w = U^S + U^B$ and it is always equal to $U^w = -c(\sigma) + w(\sigma)$.

1.2.2 Negotiation

After the seller makes specific investment σ at time $t = 0$, the seller and the buyer negotiate over the division $(a, w(\sigma) - a)$, $a \in [0, w(\sigma)]$, where a is the outcome of the negotiation process. To be specific about the negotiation process and to capture the asymmetry in parties' outside options due to investment specificity, we employ Rubinstein's alternating offers bargaining game with a one-sided outside option (see Osborne and Rubinstein (1990) and Binmore et. al. (1989)).

The game proceeds as follows. After the seller makes investment σ, the buyer and the seller learn the realization of the buyer's outside option $\underline{v}$. After learning $\underline{v}$, in period $t = 1$ the seller proposes a division $(a_1, w(\sigma) - a_1)$ to the buyer, where a_1 is the seller's share of the trade surplus and $w(\sigma) - a_1$ is the buyer's share of the trade surplus[5]. The buyer can either take the outside option $\underline{v}$, accept the seller's offer, or reject the offer and continue negotiating. If the buyer takes $\underline{v}$, the seller gets nothing, and the game ends. If the buyer accepts the division $(a_1, w(\sigma) - a_1)$, the seller produces the tradable good, sells it to the buyer, splits the trade surplus according to the negotiated division and the game ends. If the buyer rejects the division $(a_1, w(\sigma) - a_1)$, the game moves to period $t = 2$, the value of trade surplus $w(\sigma)$ depreciates by a factor of δ_1, the value of the outside option $\underline{v}$ depreciates by a factor of δ_2, and the buyer proposes the division $(a_2, \delta_1 w(\sigma) - a_2)$ to the seller where a_2 is the seller's share of the trade surplus and $\delta_1 w(\sigma) - a_2$ is the buyer's share of the trade surplus.

In period $t = 2$ the seller can either accept the buyer's offer of $(a_2, \delta_1 w(\sigma) - a_2)$ or reject it. If the seller accepts the division $(a_2, \delta_1 w(\sigma) - a_2)$, the seller produces the tradable good and sells it to the buyer, the parties receive payoffs according to the negotiated division, and the game ends. If the seller rejects the division $(a_2, \delta_1 w(\sigma) - a_2)$, the game moves to period $t = 3$, which is strategically identical to period $t = 1$, where the seller acts as a

[5]For the simplicity of the resulting equations and the ease of exposition we assume that the seller is the first to propose a division. All results hold for the case where the buyer is the first to propose.

proposer. The game continues until either any player accepts an offered division or the buyer takes the outside option. Figure 1.2 depicts the negotiation process.

It is important to emphasize three crucial elements of this negotiation process. First, only the buyer can unilaterally terminate the negotiation process by taking an outside option. The seller never unilaterally terminates the negotiation process, because, due to investment specificity, the seller's outside option always equals zero. Second, the negotiation process is costly, and the cost of the negotiation is equal to the depreciation rate δ_1 on the trade surplus and the depreciation rate δ_2 on the outside option per each negotiation period. Third, we assume a differentiated cost of negotiation in the model: the depreciation rate on the trade surplus δ_1 in general is not equal to the depreciation rate on the outside option δ_2. The distinction between the depreciation rates is necessary to capture the idea that the tradable good with specific investment (or technology) depreciates more than the same good produced with a generic investment (or technology), because a good with specific investment becomes obsolete at a faster rate than its alternative produced with generic investment. Alternatively, the lower depreciation rate on the buyer's outside option may reflect the existence of a market for used or outdated alternative goods produced with generic technology and the absence of such a market for the good produced with specific investment.

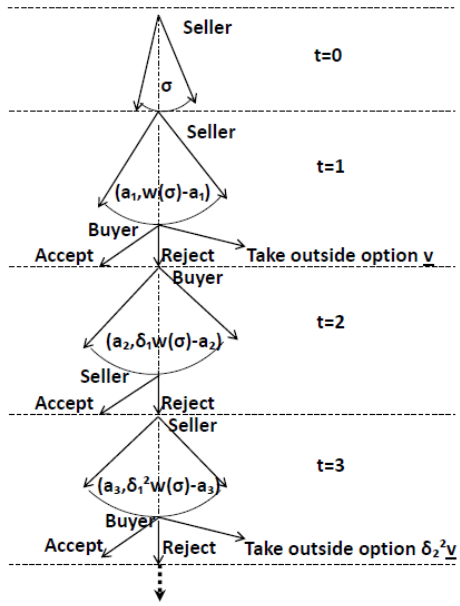

Figure 1.2. Negotiation Process

In the next proposition, we present equilibrium outcomes of the alternating offers bargaining game with the buyer's realized outside option $\underline{v}$. Following Rubinstein (1982), we use the concept of Subgame Perfect Nash Equilibrium (SPNE) to derive the equilibrium outcomes.

Proposition 1.1. *Let the negotiation process be defined as above, where $\underline{v} \in [0, V]$ is the realization of the outside option, $\delta_1 \in (0, 1)$ is the depreciation rate on the trade surplus, $\delta_2 \in (0, 1)$ is the depreciation rate on the outside option, $\underline{\delta}_1 = \frac{1}{2} - \sqrt{\frac{1}{4} - \frac{(1-\delta_2)\underline{v}}{w(\sigma)-\underline{v}}}$ and $\overline{\delta}_1 = \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{(1-\delta_2)\underline{v}}{w(\sigma)-\underline{v}}}$ are cutoff values on δ_1, and $\underline{\delta}_2 = [1 - \delta_1(1-\delta_1)(\frac{w(\sigma)}{\underline{v}} - 1)]$ is the cutoff value on δ_2. Let $x = (x_1, x_2)$ be the division proposed by the seller, where x_1 is the seller's share of the trade surplus in the seller's proposal, and let $y = (y_1, y_2)$ be the division proposed by the buyer, where y_1 is the seller's share of the trade surplus in the buyer's proposal.*

1. *If $\underline{v} \in (0, \frac{\delta_1 w(\sigma)}{1+\delta_1})$, $\delta_1 \in (0, 1)$ and $\delta_2 \in (0, 1)$, then the seller's unique SPNE strategy is*

to offer a division $x = (\frac{w(\sigma)}{1+\delta_1}, \frac{\delta_1 w(\sigma)}{1+\delta_1})$ *and to accept any division* y *with* $y_1 \geq \frac{\delta_1 w(\sigma)}{1+\delta_1}$. *The buyer's unique SPNE strategy is to offer a division* $y = (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{w(\sigma)}{1+\delta_1})$, *to accept any division* x *with* $x_2 \geq \frac{\delta_1 w(\sigma)}{1+\delta_1}$, *and never take outside option* $\underline{v}$. *The unique SPNE outcome is that* $q = 1$ *and the seller's offer* $x = (\frac{w(\sigma)}{1+\delta_1}, \frac{\delta_1 w(\sigma)}{1+\delta_1})$ *is accepted in the first period.*

2a. If $\underline{v} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, w(\sigma))$, $\delta_1 \in (0,1)$, *and* $\delta_2 \in (0, \underline{\delta}_2)$, *and if* $\underline{v} \in (\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2}, w(\sigma))$, $\delta_1 \in (0,1)$, *and* $\delta_2 \in (\underline{\delta}_2, 1)$ *then the seller's unique SPNE startegy is to offer a division* $x = (w(\sigma) - \underline{v}, \underline{v})$ *and to accept any division* y *with* $y_1 \geq \delta_1(w(\sigma) - \underline{v})$. *The buyer's unique SPNE strategy is to offer a division* $y = (\delta_1(w(\sigma) - \underline{v}), w(\sigma) - \delta_1(w(\sigma) - \underline{v}))$, *to accept any division* x *with* $x_2 \geq \underline{v}$, *and take outside option* $\underline{v}$ *if* $x_2 < \underline{v}$. *The unique SPNE outcome is that* $q = 1$ *and the seller's offer* $x = (w(\sigma) - \underline{v}, \underline{v})$ *is accepted in the first period.*

2b. If $\underline{v} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})$, $\delta_1 \in (\underline{\delta}_1, \overline{\delta}_1)$, *and* $\delta_2 \in (\underline{\delta}_2, 1)$, *then the seller's unique SPNE startegy is to offer a division* $x = (\frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}, \frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\underline{v}(\delta_2-\delta_1^2)}{1-\delta_1^2})$ *and to accept any division* y *with* $y_1 \geq \frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\delta_1\underline{v}(1-\delta_2)}{1-\delta_1^2}$. *The buyer's unique SPNE strategy is to offer a division* $y = (\frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\delta_1\underline{v}(1-\delta_2)}{1-\delta_1^2}, \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_1-\delta_1^2+\delta_1\delta_2)}{1-\delta_1^2})$, *to accept any division* x *with* $x_2 \geq \frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\underline{v}(\delta_2-\delta_1^2)}{1-\delta_1^2}$, *and never take outside option* $\underline{v}$. *The unique SPNE outcome is that* $q = 1$ *and the seller's offer* $x = (\frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}, \frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\underline{v}(\delta_2-\delta_1^2)}{1-\delta_1^2})$ *is accepted in the first period.*

3. If $w(\sigma) < \underline{v} < V$, $\delta_1 \in (0,1)$, *and* $\delta_2 \in (0,1)$, *the seller's SPNE strategy is to offer any division* x, *and the buyer's SPNE strategy is to take the outside option* $\underline{v}$. *The unique SPNE outcome is that* $q = 0$ *and there is no trade.*

Proof: see appendix.

Proposition 1.1 defines all possible negotiation outcomes as a function of different realizations of the random outside option $\underline{V}$. When the realized outside option is less than the buyer's payoff in a bargaining game without the buyer's outside option, the buyer's threat of taking the realized outside option $\underline{v}$ loses its credibility, and the buyer has to accept the SPNE offer arising in a game without the buyer's outside option. In particular, when $\underline{v} < \frac{\delta_1 w(\sigma)}{1+\delta_1}$, where $\frac{\delta_1 w(\sigma)}{1+\delta_1}$ is the buyer's payoff in the game without the buyer's outside offer, the buyer immediately agrees on the offer $(\frac{w(\sigma)}{1+\delta_1}, \frac{\delta_1 w(\sigma)}{1+\delta_1})$, which yields the seller a payoff of $\frac{w(\sigma)}{1+\delta_1} - c(\sigma)$ and the buyer a payoff of $\frac{\delta_1 w(\sigma)}{1+\delta_1}$. In this case the equilibrium outcome is unique and the trade always happens.

When the buyer's realized outside option $\underline{v}$ is larger than the buyer's payoff in a game without outside option but less than the total trade surplus, i.e. when $\underline{v} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, w(\sigma))$, the buyer is able to negotiate a payoff of at least her outside option by credibly threatening to take the realized outside option $\underline{v}$ and terminating the relationship.

The size of the buyer's additional payoff above $\underline{v}$ depends on depreciation rates δ_1 and δ_2. When δ_2 is below the cutoff value $\underline{\delta}_2$ and/or δ_1 is outside the interval $(\underline{\delta}_1, \overline{\delta}_1)$, the buyer obtains no payoff above $\underline{v}$, because the buyer's threat of rejecting the seller's offer $x = (w(\sigma) - \underline{v}, \underline{v})$ is not credible for any $\underline{v} \in (\frac{\delta_1}{1+\delta_1} w(\sigma), w(\sigma))$. As a result, the seller offers $x = (w(\sigma) - \underline{v}, \underline{v})$ and the buyer immediately accept this offer. The seller follows the same strategy independently of the values of depreciation rates δ_1 and δ_2 when the buyer's outside option is large enough, $\underline{v} \in (\frac{(\delta_1 - \delta_1^2) w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2}, w(\sigma))$. In both cases the seller's equilibrium payoff is $w(\sigma) - \underline{v} - c(\sigma)$ and the buyer's equilibrium payoff is the buyer's outside option $\underline{v}$. The equilibrium strategies in these two cases coincide with the equilibrium strategies proposed by Binmore at. al. (1989), who study an alternating-offers bargaining game with an outside option and identical discount factors on the trade surplus and the outside option (the case of $\delta = \delta_1 = \delta_2$ in our formulation).

When $\delta_1 \in (\underline{\delta}_1, \overline{\delta}_1)$, $\delta_2 > \underline{\delta}_2$ and $\underline{v} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1 - \delta_1^2) w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})$ as it is in the case 2b of Proposition 1.1, the buyer can obtain a payoff above her outside option $\underline{v}$. In this case, the seller's strategy is to offer a division $x = (\frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}, \frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\underline{v}(\delta_2-\delta_1^2)}{1-\delta_1^2})$ and the buyer's strategy is to immediately accept it. As a result, the buyer obtains an additional payoff of $\frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} - \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}$ above $\underline{v}$. The seller's payoff from the relationship is $\frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2} - c(\sigma)$ and the buyer's payoff from the relationship is equal to $\frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\underline{v}(\delta_2-\delta_1^2)}{1-\delta_1^2}$.

The size of the additional payoff $\frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} - \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}$ depends on the sizes of $\underline{v}$, δ_1 and δ_2. When the outside option does not depreciate too much and $\delta_2 > \underline{\delta}_2$ while the tradable good with specific investment depreciates at the sufficiently high rate $\delta_1 \in (\underline{\delta}_1, \overline{\delta}_1)$, the buyer can use her credible threat to reject the seller's offer and continue negotiating to extract an additional payoff above $\underline{v}$. When the outside option depreciates too much ($\delta_2 < \underline{\delta}_2$) or the tradable good does not depreciate at a high enough rate ($\delta_1 > \overline{\delta}_1$), the buyer's threat of rejecting the seller's offer loses credibility and the buyer settles on accepting a lower offer with her share equal to her outside option.

The limiting case when the outside option does not depreciate at all and δ_2 approaches 1 is particularly interesting, because the negative term $-\frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}$ in the additional payoff goes to zero and the restriction on the depreciation rate δ_1 disappears. In this case the equilibrium startegies converge to $\lim\limits_{\delta_2 \to 1}(x) = (\frac{w(\sigma)-\underline{v}}{1+\delta_1}, \underline{v} + \frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1})$ and $\lim\limits_{\delta_2 \to 1}(y) = (\frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1}, \underline{v} + \frac{w(\sigma)-\underline{v}}{1+\delta_1})$ and take a particularly simple form: the buyer obtains a guaranteed payoff of $\underline{v}$, and the parties bargain over the division of the trade surplus net of the buyer's outside option, $w(\sigma) - \underline{v}$. The outcome of the bargaining is that independently of the size of δ_1 the seller and the buyer split the trade surplus in excess of the buyer's outside, $w(\sigma) - \underline{v}$, according to the alternating-offers bargaining game without an outside option.

When the realized outside option $\underline{v}$ exceeds $w(\sigma)$, the seller cannot prevent the buyer from taking the outside option and terminating the relationship, since in this range of $\underline{v}$ the buyer strictly prefers to take the outside option to any seller's offer satisfying the seller's participation constraint. Hence, in this case the seller's payoff is $-c(\sigma)$ and the buyer's payoff is $\underline{v}$.

The figure below illustrates the seller's and the buyer's payoffs as a function of different realizations of $\underline{v}$ when $\delta_1 \in (0,1)$ and $\delta_2 < \underline{\delta}_2$, and when $\delta_1 \in (\underline{\delta}_1, \overline{\delta}_1)$ and $\underline{\delta}_2 < \delta_2$.

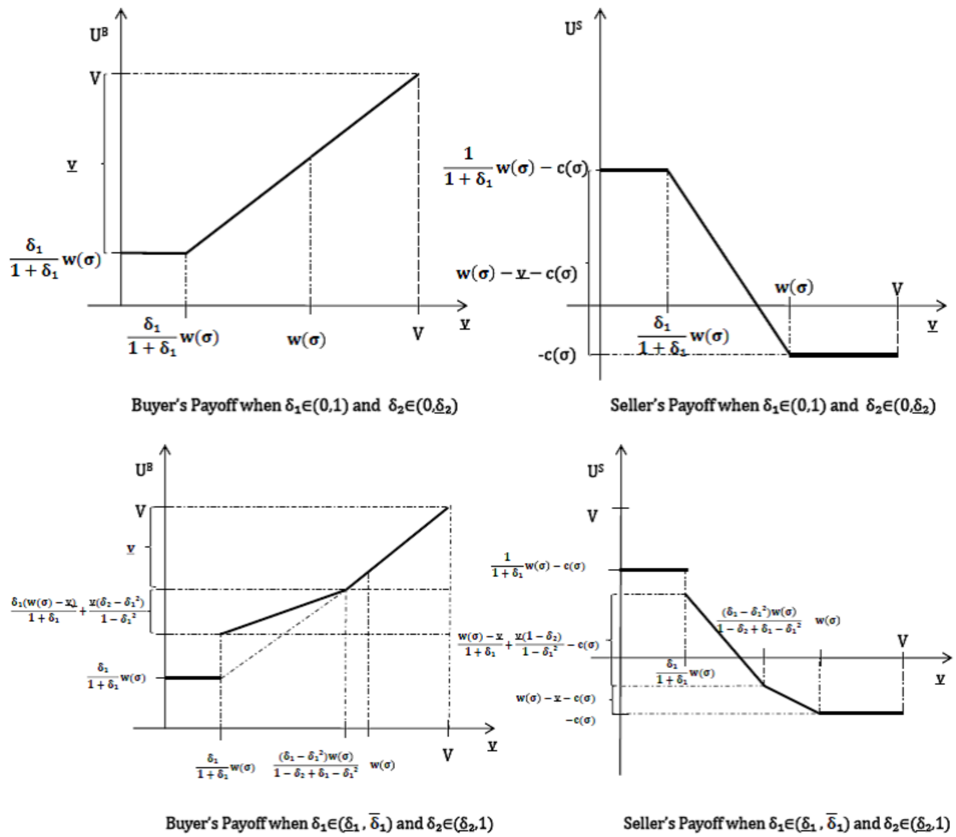

Figure 1.3. Equilibrium Payoffs of the Seller and the Buyer

The top two graphs in Figure 1.3 plot the buyer's and the seller's payoffs when the depreciation rate on the buyer's outside option is below the cutoff value $\underline{\delta}_2$ and the parties follow their equilibrium strategies. The parties' payoffs in this case are defined by parts 1, 2a, and 3 of Proposition 1.1. In particular, the buyer's payoff is equal to her outside offer $\underline{v}$ for any $\underline{v} \geq \frac{\delta_1}{1+\delta_1} w(\sigma)$ and to $\frac{\delta_1}{1+\delta_1} w(\sigma)$ when $\underline{v} < \frac{\delta_1}{1+\delta_1} w(\sigma)$.

The bottom two graphs plot the buyer's and the seller's payoffs when the depreciation rate on the buyer's outside option is above the cutoff value $\underline{\delta}_2$ and $\delta_1 \in (\underline{\delta}_1, \overline{\delta}_1)$. In this case, the buyer is able to extract an additional payoff of size $\frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} - \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}$ above $\underline{v}$ when

$\frac{\delta_1}{1+\delta_1}w(\sigma) < \underline{v} < \frac{\delta_1-\delta_1^2}{1-\delta_2+\delta_1-\delta_1^2}w(\sigma)$. For all other values of $\underline{v}$, the buyer's payoff is equal to the buyer's payoff when $\delta_2 < \underline{\delta}_2$.

1.2.3 Probabilities of Negotiation Outcomes

We assume that the buyer and the seller are risk neutral, and the parties care only about their expected payoffs. To obtain expected negotiation outcomes, we introduce a cumulative distribution function on the outside option $F_{\underline{V}}(.)$ with support $[0, V]$. Given the negotiation payoffs from Proposition 1.1, we obtain the following expected negotiation payoffs: 1) when $\underline{V} \in (0, \frac{\delta_1 w(\sigma)}{1+\delta_1})$, $\delta_1 \in (0,1)$ and $\delta_2 \in (0,1)$, the parties expect to obtain $(\frac{w(\sigma)}{1+\delta_1}, \frac{\delta_1 w(\sigma)}{1+\delta_1})$, 2a) when $\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, w(\sigma))$, $\delta_1 \in (0,1)$ and $\delta_2 \in (0, \underline{\delta}_2^{EV})$, the parties expect to obtain $(w(\sigma) - E[\underline{V}|\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, w(\sigma))], E[\underline{V}|\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, w(\sigma))])$, and when $\underline{V} \in (\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2}, w(\sigma))$, $\delta_1 \in (0,1)$ and $\delta_2 \in (\underline{\delta}_2^{EV}, 1)$, the parties expect to obtain $(w(\sigma) - E[\underline{V}|\underline{V} \in (\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2}, w(\sigma))], E[\underline{V}|\underline{V} \in (\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2}, w(\sigma))])$, 2b) when $\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})$, $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\delta_2 \in (\underline{\delta}_2^{EV}, 1)$, the parties expect to obtain $(\frac{w(\sigma)}{1+\delta_1} - \frac{\delta_2-\delta_1}{1-\delta_1^2}E[\underline{V}|\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})]), \frac{\delta_1 w(\sigma)}{1+\delta_1} + \frac{\delta_2-\delta_1}{1-\delta_1^2}E[\underline{V}|\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})])$, and 3) when $w(\sigma) < \underline{V} < V$, $\delta_1 \in (0,1)$ and $\delta_2 \in (0,1)$, the parties expect to obtain $(0, E[\underline{V}|\underline{V} \in (w(\sigma), V)])$.

Following the introduction of expected negotiation payoffs, the cutoff values $\underline{\delta}_1$, $\overline{\delta}_1$, and $\underline{\delta}_2$ have to be accordingly modified. We introduce ex ante cutoff values on δ_1, $\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV} = \frac{1}{2} \mp \sqrt{\frac{1}{4} - \frac{(1-\delta_2)E[\underline{V}|\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})]}{w(\sigma) - E[\underline{V}|\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})]}}$ an ex ante cutoff value on δ_1, $\underline{\delta}_2^{EV} = [1 - \delta_1(1-\delta_1)(\frac{w(\sigma)}{E[\underline{V}|\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})]} - 1)]$.

With cdf $F_{\underline{V}}(.)$ we can define probabilities of four negotiation outcomes in the case when $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\delta_2 \in (\underline{\delta}_2^{EV}, 1)$. Let $\lambda_1(\sigma) = F_{\underline{V}}(0 < \underline{V} < \frac{\delta_1 w(\sigma)}{1+\delta_1})$, $\lambda_2(\sigma) = F_{\underline{V}}(\frac{\delta_1 w(\sigma)}{1+\delta_1} < \underline{V} < \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})$, $\lambda_3(\sigma) = F_{\underline{V}}(\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2} < \underline{V} < w(\sigma))$ and $\lambda_4(\sigma) = F_{\underline{V}}(w(\sigma) < \underline{V} < V)$. In the case when $\delta_1 \in (0,1)$ and $\delta_2 \in (0, \underline{\delta}_2^{EV})$, there are only three probabilities,

$\lambda_1(\sigma) = F_{\underline{V}}(0 < \underline{V} < \frac{\delta_1 w(\sigma)}{1+\delta_1})$, $1 - \lambda_1(\sigma) - \lambda_4(\sigma) = F_{\underline{V}}(\frac{\delta_1 w(\sigma)}{1+\delta_1} < \underline{V} < w(\sigma))$, and $\lambda_4(\sigma) = F_{\underline{V}}(w(\sigma) < \underline{V} < V)$.

The set of probabilities $\langle \lambda_1(\sigma), \lambda_2(\sigma), \lambda_3(\sigma), \lambda_4(\sigma) \rangle$ determines an ex ante negotiation position of the seller vis-a-vis the buyer. A set of probabilities, where $\lambda_1(\sigma)$ is small in comparison to $1 - \lambda_1(\sigma)$ corresponds to a weaker ex ante negotiation position of the seller relative to the buyer. When $\lambda_1(\sigma)$ is small, the seller is more likely to concede a higher share of the trade surplus to prevent the buyer from taking the outside option or is more likely to face the termination of the relationship. Respectively, a set of probabilities, where $1 - \lambda_1(\sigma)$ are small in comparison to $\lambda_1(\sigma)$, corresponds to a stronger ex ante negotiation position of the seller relative to the buyer, because with a high $\lambda_1(\sigma)$ the seller is more likely to face the most favorable negotiation outcome.

The set of probabilities $\langle \lambda_1(\sigma), \lambda_2(\sigma), \lambda_3(\sigma), \lambda_4(\sigma) \rangle$ depends on the cdf of the outside option $F_{\underline{V}}(.)$ and the value of the trade surplus $w(\sigma)$. Since $F_{\underline{V}}(.)$ is independent of σ, we can determine how $\lambda_1(\sigma)$, $\lambda_2(\sigma)$, $\lambda_3(\sigma)$, and $\lambda_4(\sigma)$ depend on σ. In the next proposition we establish a relationship between specific investment σ and the set $\langle \lambda_1(\sigma), \lambda_2(\sigma), \lambda_3(\sigma), \lambda_4(\sigma) \rangle$.

Proposition 1.2.

For any distribution function $F_{\underline{V}}(.)$, an increase in specific investment σ raises the probability of trade $\lambda_1(\sigma)$, $\frac{\partial \lambda_1(\sigma)}{\partial \sigma} > 0$, and lowers the probability of termination of the relationship $\lambda_4(\sigma)$, $\frac{\partial \lambda_4(\sigma)}{\partial \sigma} < 0$. The effect of specific investment on $\lambda_2(\sigma)$ depends on the relative effects of specific investment on $\lambda_1(\sigma)$, $\lambda_3(\sigma)$ and $\lambda_4(\sigma)$. In particular, if $-\frac{\partial \lambda_3(\sigma)}{\partial \sigma} - \frac{\partial \lambda_4(\sigma)}{\partial \sigma} > \frac{\partial \lambda_1(\sigma)}{\partial \sigma}$, then $\frac{\partial \lambda_2(\sigma)}{\partial \sigma} > 0$, otherwise $\frac{\partial \lambda_2(\sigma)}{\partial \sigma} < 0$. The effect of specific investment on $\lambda_3(\sigma)$ depends on the relative effects of specific investment on $\lambda_1(\sigma)$, $\lambda_2(\sigma)$, and $\lambda_3(\sigma)$. In particular, $-\frac{\partial \lambda_2(\sigma)}{\partial \sigma} - \frac{\partial \lambda_4(\sigma)}{\partial \sigma} > \frac{\partial \lambda_1(\sigma)}{\partial \sigma}$, then $\frac{\partial \lambda_3(\sigma)}{\partial \sigma} > 0$, otherwise $\frac{\partial \lambda_3(\sigma)}{\partial \sigma} < 0$.

Proof: By definition, $\lambda_1(\sigma) = F_{\underline{V}}(0 < \underline{V} < \frac{\delta_1 w(\sigma)}{1+\delta_1})$. Consider an increase in specific investment from σ to σ'. By assumption, an increase in specific investment from σ to σ' raises the buyer's valuation of the tradable good from $b(\sigma)$ to $b(\sigma')$ and lowers the seller's cost of the tradable good from $s(\sigma)$ to $s(\sigma')$. This means that $w(\sigma) = b(\sigma) - s(\sigma) <$

$b(\sigma') - s(\sigma') = w(\sigma')$.

Since the distribution function of the outside option $F_{\underline{V}}(.)$ is independent of specific investment and the trade surplus $w(\sigma)$ increases with investment, the probability of trade $\lambda_1(\sigma)$ unconditionally increases in specific investment,

$$\lambda_1(\sigma) = F_{\underline{V}}(0 < \underline{V} < \frac{\delta_1 w(\sigma)}{1+\delta_1}) < F_{\underline{V}}(0 < \underline{V} < \frac{\delta_1 w(\sigma')}{1+\delta_1}) = \lambda_1(\sigma').$$

Thus, we conclude that the probability of trade $\lambda_1(\sigma)$ increases in specific investment.

By the same argument we can show that an increase in specific investment lowers the probability of termination $\lambda_4(\sigma)$. At the same time we cannot determine how the probabilities of trade $\lambda_2(\sigma)$ and $\lambda_3(\sigma)$ respond to an increase in investment σ, since investment effects on $\lambda_2(\sigma)$ and $\lambda_3(\sigma)$ depend on the shape of a particular distribution function $F_{\underline{V}}(.)$ and the value of $w(\sigma)$. However, since $\lambda_1(\sigma)$, $\lambda_2(\sigma)$, $\lambda_3(\sigma)$, and $\lambda_4(\sigma)$ form a partition, $\sum_{i=1}^{4} \lambda_i(\sigma) = 1$ and $\lambda_2(\sigma) = 1 - \lambda_1(\sigma) - \lambda_3(\sigma) - \lambda_4(\sigma)$, then $\frac{\partial \lambda_2(\sigma)}{\partial \sigma} > 0$ if and only if $-\frac{\partial \lambda_3(\sigma)}{\partial \sigma} - \frac{\partial \lambda_4(\sigma)}{\partial \sigma} > \frac{\partial \lambda_1(\sigma)}{\partial \sigma}$, while $\frac{\partial \lambda_2(\sigma)}{\partial \sigma} < 0$ if and only if $-\frac{\partial \lambda_3(\sigma)}{\partial \sigma} - \frac{\partial \lambda_4(\sigma)}{\partial \sigma} > \frac{\partial \lambda_1(\sigma)}{\partial \sigma}$. The same argument applies to $\lambda_3(\sigma)$.

The resolution of the hold-up problem in this chapter critically depends on the effect of specific investment on the probabilities $\lambda_1(\sigma)$, $\lambda_2(\sigma)$, $\lambda_3(\sigma)$ and $\lambda_4(\sigma)$. This is in contrast with prior approaches to the hold-up problem, which assume that specific investment has no effect on the probabilities of negotiation outcomes, i.e. that $\lambda_1'(\sigma) = \lambda_2'(\sigma) = \lambda_3'(\sigma) = \lambda_4'(\sigma) = 0$.

1.2.4 Technical Assumptions

1.1We assume that specific investment σ raises the cost of investment $c(\sigma)$, $c'(\sigma) > 0$, lowers the seller's cost of production of the tradable good $s(\sigma)$, $s'(\sigma) < 0$, and raises the buyer's valuation of the tradable good $b(\sigma)$, $b'(\sigma) > 0$. This implies that the trade surplus $w(\sigma) = b(\sigma) - s(\sigma)$ also increases in investment, $w'(\sigma) > 0$.

1.2 We further assume that the trade surplus $w(\sigma)$ is concave, $w''(\sigma) < 0$, and the seller's cost of investment $c(\sigma)$ is convex, $c''(\sigma) > 0$.

1.3 We impose two second-order concavity conditions on the cdf $F_{\underline{V}}(.)$ and functions $w(\sigma)$ and $c(\sigma)$. These two conditions together with the concavity requirement of $w(\sigma)$ and the convexity requirement of $c(\sigma)$ guarantee that the joint and the seller's objective functions are globally concave:

$(A1.3.1)$ $-\lambda_4'(\sigma)w'(\sigma) < c''(\sigma) - w''(\sigma)(1 - \lambda_4(\sigma))$,

$(A1.3.2)$ $w'(\sigma)(\frac{\lambda_1'(\sigma)(1-\delta_1^2+\delta_2-\delta_1)}{(1-\delta_1)(1+\delta_1)^2} + \frac{\lambda_2'(\sigma)}{1+\delta_1} + \lambda_3'(\sigma)) + \frac{\lambda_1''(\sigma)w(\sigma)(\delta_2-\delta_1)}{(1-\delta_1)(1+\delta_1)^2} <$
$c''(\sigma) - w''(\sigma)(\frac{\lambda_1(\sigma)+\lambda_2(\sigma)}{1+\delta_1} + \lambda_3(\sigma))$.

The right-hand side of condition A1.3.1 is strictly positive since $c''(\sigma) > 0$ and $w''(\sigma) < 0$, while the right-hand side is strictly positive by Proposition 1.2, where we show that $\lambda_4'(\sigma) < 0$. Similarly, the right-hand side of condition A1.3.2 is strictly positive because $c''(\sigma) > 0$ and $w''(\sigma) < 0$, while the sign of the right-hand side is indeterminate and depends on the sign of $\lambda_1''(\sigma)$ and the magnitude and signs of $\lambda_2'(\sigma) + \lambda_3'(\sigma)$. We do not impose any restrictions on the distribution function $F_{\underline{V}}(.)$ beyond those implied by the second-order concavity conditions A1.3.1 and A1.3.2.

1.4. We assume that $b(\sigma) > s(\sigma)$, so that the trade is always efficient.

1.5. We assume that for $\sigma \in [0, \Sigma]$, $w(\sigma) > c(\sigma)$, so that it is socially efficient to undertake specific investment.

1.3 First-Best Specific Investment

In the first-best formulation, a social planner maximizes the joint expected payoff to the buyer and the seller. The expectation is defined over all possible realizations of the outside option for any $\delta_1 \in (0,1)$ and $\delta_2 \in (0,1)$. In particular, if the outside option $\underline{V}$ is in the interval $[0, w(\sigma)]$, the trade occurs, and the joint payoff is equal to the expected trade surplus $E[w(\sigma)|\underline{V} \in [0, w(\sigma)]]$. Since the outside option $\underline{V}$ and specific investment σ are independent, $E[w(\sigma)|\underline{V} \in [0, w(\sigma)]] = w(\sigma)$.

If the outside option $\underline{V}$ is in the interval $(w(\sigma), V]$, the trade doesn't occur, and the joint payoff is equal to the buyer's expected outside option conditional on exceeding the total trade surplus, $E[\underline{V}|\underline{V} \in (w(\sigma), V]]$.

The trade doesn't occur with probability $\lambda_4(\sigma)$, and the trade occurs with probability $1 - \lambda_4(\sigma) = \lambda_1(\sigma) + \lambda_2(\sigma) + \lambda_3(\sigma)$. In addition, whether the trade occurs or not, the seller undertakes specific investment σ at cost $c(\sigma)$. The joint expected social payoff to the buyer and the seller EU^W is given in equation (1.1).

(1.1) $EU^W = -c(\sigma) + (1 - \lambda_4(\sigma))w(\sigma) + \lambda_4(\sigma)E[\underline{V}|\underline{V} \in (w(\sigma), V)]$

Given that the equation (1.1) is globally concave by the second-order concavity condition A1.3.1, the first-best level of specific investment $\sigma_{fb} = \underset{\sigma \geq 0}{argmax} EU^W$ is a maximizer of EU^W. As a result, the first-best level of specific investment σ_{fb} should satisfy the the first-order maximization equation (1.2).

(1.2) $(1 - \lambda_4(\sigma_{fb}))w'(\sigma_{fb}) + \lambda_4(\sigma_{fb})\frac{\partial E[\underline{V}|\underline{V} \in (w(\sigma), V)]}{\partial \sigma_{fb}} =$

$c'(\sigma_{fb}) - \lambda_4'(\sigma_{fb})(E[\underline{V}|\underline{V} \in (w(\sigma_{fb}), V)] - w(\sigma_{fb}))$

Equation (1.2) has a natural interpretation. The left-hand side of equation (1.2) is the expected marginal benefit from the first-best specific investment σ_{fb}, while the right-hand side is the expected marginal cost of σ_{fb}. The expected marginal benefit from σ_{fb} consists of two positive components. The first component is the expected increase in the trade surplus shared by the buyer and the seller. This component is positive given assumption 1.1. The second component is the expected increase in the buyer's expected outside option conditional on exceeding the total trade surplus. The second component is strictly positive as well. The exact derivation of the term $\frac{\partial E[\underline{V}|\underline{V} \in (w(\sigma), V)]}{\partial \sigma_{fb}}$ and its sign are presented in the appendix.

The right-hand side of equation (1.2) defines the expected marginal cost of specific investment, and it consists of two components as well: an increase in the cost of sunk specific investment $c'(\sigma_{fb})$ shared only by the seller and the marginal cost of adjustment in the ex ante negotiation position of the seller vis-a-vis the buyer. The extent of the marginal cost from adjustment in the ex ante negotiation position depends on the size of the strictly negative $\lambda_4'(\sigma_{fb})$ and the positive difference between $E[\underline{V}|\underline{V} \in (w(\sigma_{fb}), V)]$ and $w(\sigma_{fb})$.

Equation (1.2) can be simplified if we expand the term $\frac{\partial E[\underline{V}|\underline{V} \in (w(\sigma),V)]}{\partial \sigma_{fb}}$ to get the first-best optimality condition (1.3):

$$(1.3)\ \ (1 - \lambda_4(\sigma_{fb}))w'(\sigma_{fb}) = c'(\sigma_{fb}).$$

Similar to equation (1.2), in equation (1.3) the left-hand side denotes the expected marginal benefit consisting of the jointly shared expected increase in the trade surplus $w(\sigma_{fb})$. The right-hand side denotes the expected marginal cost of specific investment $c'(\sigma_{fb})$ shared only by the seller. In the optimum the expected marginal benefit from specific investment shared by both parties should be equal to the marginal cost of specific investment imposed only on the seller.

1.4 Equilibrium Specific Investment

In equilibrium the seller chooses a level of specific investment to maximize her expected benefit from the relationship with the buyer. The seller's expected benefit from the relationship EU^S is a sum of the cost of specific investment $c(\sigma)$ and the expected negotiation payoffs. Since the expected negotiation payoffs depend on the values of δ_1 and δ_2, we consider two cases: when (a) $\delta_1 \in (0,1)$ and $\delta_2 < \underline{\delta}_2^{EV}$ and (b) $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\underline{\delta}_2^{EV} < \delta_2$. Equation (1.4a) presents the seller's expected payoff when $\delta_1 \in (0,1)$ and $\delta_2 < \underline{\delta}_2^{EV}$, and equation (1.4b) presents the seller's expected payoff when $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\underline{\delta}_2^{EV} < \delta_2$.

$$(1.4a)\ \ EU_a^S = \lambda_1(\sigma)E[\tfrac{w(\sigma)}{1+\delta_1}|\underline{V} \in (0, \tfrac{\delta_1 w(\sigma)}{1+\delta_1})]) +$$
$$(1 - \lambda_1(\sigma) - \lambda_4(\sigma))E[w(\sigma) - \underline{V}|\underline{V} \in (\tfrac{\delta_1 w(\sigma)}{1+\delta_1}, w(\sigma))]) - c(\sigma)$$

$$(1.4b)\ \ EU_b^S = \lambda_1(\sigma)E[\tfrac{w(\sigma)}{1+\delta_1}|\underline{V} \in (0, \tfrac{\delta_1 w(\sigma)}{1+\delta_1})]) +$$
$$\lambda_2(\sigma)E[\tfrac{w(\sigma)-\underline{V}}{1+\delta_1} + \tfrac{\underline{V}(1-\delta_2)}{1-\delta_1^2}|\underline{V} \in (\tfrac{\delta_1 w(\sigma)}{1+\delta_1}, \tfrac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})]) +$$
$$\lambda_3(\sigma)E[w(\sigma) - \underline{V}|\underline{V} \in (\tfrac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2}, w(\sigma))]) - c(\sigma)$$

Since $w(\sigma)$ and $\underline{V}$ are independent, and $\underline{V}$ enters linearly into the seller's equilibrium payoffs, we can simplify equations (1.4a) and (1.4b) to obtain equations (1.5a) and (1.5b):

$$(1.5a)\ \ EU_a^S = (1-\lambda_1(\sigma)-\lambda_4(\sigma))(w(\sigma)-E[\underline{V}|\underline{V}\in(\tfrac{\delta_1 w(\sigma)}{1+\delta_1},w(\sigma))])+ \\ \lambda_1(\sigma)\tfrac{w(\sigma)}{1+\delta_1}-c(\sigma).$$

$$(1.5b)\ \ EU_b^S = \lambda_1(\sigma)\tfrac{w(\sigma)}{1+\delta_1}+\lambda_3(\sigma)(w(\sigma)-E[\underline{V}|\underline{V}\in(\tfrac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2},w(\sigma))])+ \\ \lambda_2(\sigma)(\tfrac{w(\sigma)}{1+\delta_1}-\tfrac{\delta_2-\delta_1}{1-\delta_1^2}E[\underline{V}|\underline{V}\in(\tfrac{\delta_1 w(\sigma)}{1+\delta_1},\tfrac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})])-c(\sigma).$$

Given that equations (1.5a) and (1.5b) are globally concave by the second-order concavity conditions A1.3.1 and A1.3.2, the equilibrium specific investment $\sigma_{eq}^a = \underset{\sigma\geq 0}{argmax} EU_a^S$ is a maximizer of EU_a^S and the equilibrium specific investment $\sigma_{eq}^b = \underset{\sigma\geq 0}{argmax} EU_b^S$ is a maximizer of EU_b^S. Hence, the equilibrium specific investments σ_{eq}^a and σ_{eq}^b should satisfy the first-order maximization conditions (1.6a) and (1.6b).

$$(1.6a)\ \ \tfrac{\lambda_1(\sigma_{eq}^a)w'(\sigma_{eq}^a)+\lambda_1'(\sigma_{eq}^a)w(\sigma_{eq}^a)}{1+\delta_1}+(1-\lambda_1(\sigma_{eq}^a)-\lambda_4(\sigma_{eq}^a))(w'(\sigma_{eq}^a)-\tfrac{\partial E[\underline{V}|\underline{V}\in(\frac{\delta_1 w(\sigma_{eq}^a)}{1+\delta_1},w(\sigma_{eq}^a))]}{\partial\sigma_{eq}^a})+ \\ (-\lambda_1'(\sigma_{eq}^a)-\lambda_4'(\sigma_{eq}^a))(w(\sigma_{eq}^a)-E[\underline{V}\in(\tfrac{\delta_1 w(\sigma_{eq}^a)}{1+\delta_1},w(\sigma_{eq}^a))])-c'(\sigma_{eq}^a)=0$$

$$(1.6b)\ \ \tfrac{\lambda_1(\sigma_{eq}^b)w'(\sigma_{eq}^b)}{1+\delta_1}+\tfrac{\lambda_1'(\sigma_{eq}^b)w(\sigma_{eq}^b)}{1+\delta_1}+\lambda_3(\sigma_{eq}^b)(w'(\sigma_{eq}^b)-\tfrac{\partial E[\underline{V}|\underline{V}\in(\frac{(\delta_1-\delta_1^2)w(\sigma_{eq}^b)}{1-\delta_2+\delta_1-\delta_1^2},w(\sigma_{eq}^b))]}{\partial\sigma_{eq}^b})+ \\ \lambda_3'(\sigma_{eq}^b)(w(\sigma_{eq}^b)-E[\underline{V}|\underline{V}\in(\tfrac{(\delta_1-\delta_1^2)w(\sigma_{eq}^b)}{1-\delta_2+\delta_1-\delta_1^2},w(\sigma_{eq}^b))])-c'(\sigma_{eq}^b)+ \\ \lambda_2(\sigma_{eq}^b)(\tfrac{w'(\sigma_{eq}^b)}{1+\delta_1}-\tfrac{\delta_2-\delta_1}{1-\delta_1^2}\tfrac{\partial E[\underline{V}|\underline{V}\in(\frac{\delta_1 w(\sigma_{eq}^b)}{1+\delta_1},\frac{(\delta_1-\delta_1^2)w(\sigma_{eq}^b)}{1-\delta_2+\delta_1-\delta_1^2})]}{\partial\sigma_{eq}^b})+ \\ \lambda_2'(\sigma_{eq}^b)(\tfrac{w(\sigma_{eq}^b)}{1+\delta_1}-\tfrac{\delta_2-\delta_1}{1-\delta_1^2}E[\underline{V}|\underline{V}\in(\tfrac{\delta_1 w(\sigma_{eq}^b)}{1+\delta_1},\tfrac{(\delta_1-\delta_1^2)w(\sigma_{eq}^b)}{1-\delta_2+\delta_1-\delta_1^2})])=0$$

Equations (1.6a) and (1.6b) can be simplified into equations (1.7a) and (1.7b) (see chapter appendix for details).

$$(1.7a)\ \ w'(\sigma_{eq}^a)(\tfrac{\delta_1\lambda_1(\sigma_{eq}^a)}{1+\delta_1}+1-\lambda_4(\sigma_{eq}^a))=c'(\sigma_{eq}^a)$$

$$(1.7b)\ \ w'(\sigma_{eq}^b)(\tfrac{\lambda_1(\sigma_{eq}^b)+\lambda_2(\sigma_{eq}^b)}{1+\delta_1}+\lambda_3(\sigma_{eq}^b))+\lambda_1'(\sigma_{eq}^b)\tfrac{w(\sigma_{eq}^b)(\delta_2-\delta_1)}{(1-\delta_1)(1+\delta_1)^2}=c'(\sigma_{eq}^b)$$

The interpretation of equations (1.7a) and (1.7b) is similar to the interpretation of equation (1.3). The left-hand sides of equations (1.7a) and (1.7b) are the seller's expected marginal benefits from equilibrium investments σ_{eq}^a and σ_{eq}^b, and the right-hand sides are the seller's expected marginal costs from σ_{eq}^a and σ_{eq}^b. The left-hand side of equation (1.7a) is a sum of the seller's expected marginal benefits from an increase in the trade surplus in all negotiation outcomes when $\delta_1 \in (0,1)$ and $\delta_2 < \underline{\delta}_2^{EV}$. Similarly, the left-hand side of equation (1.7b) is a sum of the seller's marginal benefits from an increase in the trade surplus when $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\underline{\delta}_2^{EV} < \delta_2$. As opposed to equation (1.7a), the left-hand side of equation (1.7b) includes an additional marginal benefit term: the increase in the seller's expected marginal benefit from improvement in the seller's ex ante negotiation position. More precisely, the strictly positive term $\lambda_1'(\sigma_{eq}^b)\frac{w(\sigma_{eq}^b)(\delta_2-\delta_1)}{(1-\delta_1)(1+\delta_1)^2}$ denotes the expected marginal benefit of the seller due to improvement in the seller's negotiation position of size $\lambda_1'(\sigma_{sb}^b)$ multiplied by the associated payoff of size $\frac{w(\sigma_{eq}^b)(\delta_2-\delta_1)}{(1-\delta_1)(1+\delta_1)^2}$ (by Proposition 1.2, $\lambda_1'(\sigma_{eq}^b) > 0$). The right-hand sides of equations (1.7a) and (1.7b), or the cost sides, consist of a common single term, which is the strictly positive increase in the physical cost of investment of size $c'(\sigma_{eq}^a)$ or $c'(\sigma_{eq}^b)$.

Equations (1.7a) characterizes the seller's equilibrium specific investment σ_{eq}^a when $\delta_1 \in (0,1)$ and $\delta_2 < \underline{\delta}_2^{EV}$, while equation (1.7b) characterizes the seller's equilibrium level of specific investment when $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\underline{\delta}_2^{EV} < \delta_2$. By comparing these two equations we can draw a conclusion about whether the equilibrium specific investment is higher in case (a) or in case (b).

Since the cost function $c(.)$ is strictly convex, we conclude that the level of specific investment when $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\underline{\delta}_2^{EV} < \delta_2$ is higher than the level of specific investment when $\delta_1 \in (0,1)$ and $\delta_2 < \underline{\delta}_2^{EV}$ if and only if the left-hand side of equation (1.7b) is exceeds the left-hand side of equation (1.7a) for some common σ. We summarize this finding in the Proposition 1.3.

Proposition 1.3.

For any convex function $c(.)$, $\sigma_{eq}^b > \sigma_{eq}^a$ if and only if $\lambda_1'(\sigma_{eq}^b)w(\sigma_{eq}^b)\frac{\delta_2-\delta_1}{\delta_1(1-\delta_1^2)} > \lambda_2(\sigma_{eq}^b)w'(\sigma_{eq}^b)$, $\sigma_{eq}^b < \sigma_{eq}^a$ if and only if $\lambda_1'(\sigma_{eq}^b)w(\sigma_{eq}^b)\frac{\delta_2-\delta_1}{\delta_1(1-\delta_1^2)} < \lambda_2(\sigma_{eq}^b)w'(\sigma_{eq}^b)$, and $\sigma_{eq}^b = \sigma_{eq}^a$ if and only if $\lambda_1'(\sigma_{eq}^b)w(\sigma_{eq}^b)\frac{\delta_2-\delta_1}{\delta_1(1-\delta_1^2)} = \lambda_2(\sigma_{eq}^b)w'(\sigma_{eq}^b)$.

Proposition 1.3 tells that the equilibrium level of investment σ_{eq}^b (when $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\underline{\delta}_2^{EV} < \delta_2$) does not unconditionally exceed the equilibrium level of investment σ_{eq}^a

(when $\delta_1 \in (0,1)$ and $\underline{\delta}_2^{EV} < \delta_2$). Whether σ_{eq}^b exceeds σ_{eq}^a or not, depends on the distribution function of the outside option, the trade surplus function $w(.)$, and depreciation rates δ_1 and δ_2.

Nevertheless, we can state that $\sigma_{eq}^b < \sigma_{eq}^a$ with certainty in two cases: when $\delta_1 \to \delta_2$ and when $\lambda_1'(\sigma_{eq}^b) \to 0$. The case when $\delta_1 \to \delta_2$ reflects the situation when the depreciation rate on the tradable good with specific investment converges to the depreciation rate on the outside option with generic investment. The case when $\lambda_1'(\sigma_{eq}^b) \to 0$ defines the situation when specific investment has no effect on the likelihood of negotiation outcomes, which is possible if the increase in the trade surplus from specific investment is matched by an equal sized increase in the value of the buyer's outside option. Both cases contradict the two main assumptions of specific investment: the absence of the effect of specific investment on the outside option and the higher depreciation rate on the good produced with specific investment.

1.5 First-Best Optimality of Equilibrium Investment

The equilibrium levels of specific investments σ_{eq}^a and σ_{eq}^b are socially optimal, if they also satisfy the first-best optimality condition (1.3). The first-best optimality condition (1.3) can be rearranged into a sum of the seller's and the buyer's expected marginal payoffs from specific investment. By subtracting equations (1.7a) and (1.7b) from the rearranged equation (1.3), we obtain equations (1.8a) and (1.8b).

$(1.8a)\ \ \lambda_1(\sigma_{eq}^a)\frac{\delta_1 w'(\sigma_{eq}^a)}{1+\delta_1} = 0$

$(1.8b)\ \ (\lambda_1(\sigma_{sb}^b) + \lambda_2(\sigma_{eq}^b))\frac{\delta_1 w'(\sigma_{eq}^a)}{1+\delta_1} = \lambda_1'(\sigma_{eq}^b)\frac{w(\sigma_{eq}^b)(\delta_2-\delta_1)}{(1-\delta_1)(1+\delta_1)^2}$

Equation (1.8a) is the social optimality condition for the equilibrium specific investment σ_{eq}^a when $\delta_1 \in (0,1)$ and $\delta_2 < \underline{\delta}_2^{EV}$ and equation (1.8b) is the social optimality condition for the equilibrium specific investment σ_{eq}^b when $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\underline{\delta}_2^{EV} < \delta_2$. Both equations show that for the seller's choice of specific investment to be socially optimal, the equilibrium specific investment should equate the buyer's expected marginal benefit to the buyer's expected marginal cost. The left-hand sides of equations (1.8a) and (1.8b) indicate the buyer's expected marginal benefit from the seller's investment, while the right-hand sides indicate the buyer's expected marginal cost.

Equation (1.8a) shows that when the outside option depreciates too much independently of δ_1 and $\delta_2 < \underline{\delta}_2^{EV}$, the buyer does not face any costs of specific investment as indicated by the zero right-hand side of equation (1.8a) and a strictly positive left-hand side of equation (1.8a). This is because with probability λ_1 the buyer obtains a positive share of the trade surplus and with probability $1 - \lambda_1$ the buyer obtains a payoff equal to her expected outside option independently of the level of specific investment. Hence, an increase in specific investment strictly raises the buyer's payoff with probability λ_1 without any impact on the buyer's payoff in any other negotiation outcome. As a result, in case (a) the buyer's expected marginal benefit from investment always exceeds the buyer's expected marginal cost of zero, and the seller's choice of specific investment always results in under-investment.

Equation (1.8b) shows that when the depreciation rate on the outside option is sufficiently low ($\underline{\delta}_2^{EV} < \delta_2$) and the tradable good with specific investment depreciates at a high enough rate $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$, the buyer receives both the expected marginal benefit from the seller's investment and the expected marginal cost. In particular, the strictly positive left-hand side of equation (1.8b), or the marginal benefit side, is the expected marginal benefit from an increase in the trade surplus. The strictly positive right-hand side of equation (1.8b), or the marginal cost side, is the buyer's expected marginal cost from an increase in the probability $\lambda_1(\sigma_{eq}^b)$. By Proposition 1.2, the buyer's expected marginal cost is strictly positive, since $\lambda_1'(\sigma_{eq}^b) > 0$.

The buyer's expected marginal cost of $\lambda_1'(\sigma_{eq}^b)\frac{w(\sigma_{eq}^b)(\delta_2-\delta_1)}{(1-\delta_1)(1+\delta_1)^2}$ indicates the buyer's loss of the additional payoff above her outside option when the probability of the negotiation outcome $\lambda_1(\sigma)$ increases. The buyer does not lose any such payoff in case (a), because in case (a) the high depreciation rate on the outside option does not allow the buyer to crediby use her threat of rejecting the seller's offer and extract an additional payoff above her outside option.

Hence, the hold-up problem can be resolved only in case (b), when $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\underline{\delta}_2^{EV} < \delta_2$. In particular, depending on the distribution function of the outside option, the trade surplus function and the depreciation rates, it is possible to have a socially optimal investment if condition (1.8b) holds as equality, under-investment if condition (8b) is a positive inequality, and over-investment if condition (1.8b) is a negative inequality.

In the limiting case when $\delta_2 \to \underline{\delta}_2^{EV}$, equations (1.8b) converges to equation (1.8a) and the hold-up problem always persists irrespective of the depreciation rate δ_1 and functional forms of the trade surplus function $w(.)$ and the density function $F_{\underline{V}}(.)$. The next proposition precisely defines conditions for the social optimality of the seller's specific investment for all values of δ_1 and δ_2.

Proposition 1.4.

Let σ_{eq}^a and σ_{eq}^b maximize the seller's objective functions (1.5a) and (1.5b) respectively, $\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV} = \frac{1}{2} \mp \sqrt{\frac{1}{4} - \frac{(1-\delta_2)E[\underline{V}|\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})]}{w(\sigma)-E[\underline{V}|\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})]}}$ *be the ex ante cutoff values on δ_1 and* $\underline{\delta}_2^{EV} = [1 - \delta_1(1 - \delta_1)(\frac{w(\sigma)}{E[\underline{V}|\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})]} - 1)]$ *be the ex ante cutoff value on δ_2. If $\delta_1 \in (0,1)$ and $\underline{\delta}_2^{EV} > \delta_2$, the seller's equilibrium investment σ_{eq}^a always results in the under-invests from the socially optimal point of view. If $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\underline{\delta}_2^{EV} < \delta_2$, the seller's equilibrium investment σ_{eq}^b is ex ante socially optimal, if condition (1.8b) holds as equality. If $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\underline{\delta}_2^{EV} < \delta_2$ and the left-hand side of condition (1.8b) exceeds its right-hand side, the seller ex ante under-invests. If $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\underline{\delta}_2^{EV} < \delta_2$ and the left-hand side of condition (1.8b) exceeds its right-hand side, the seller ex ante over-invests.*

Proof: The proof is trivial. Since σ_{eq}^a can never satisfy equation (1.8a) as an equality, it follows that when $\underline{\delta}_2^{EV} > \delta_2$, the social optimality condition (1.3) is never satisfied and the seller always under-invests. When $\delta_1 \in (\underline{\delta}_1^{EV}, \overline{\delta}_1^{EV})$ and $\underline{\delta}_2^{EV} < \delta_2$, equilibrium investment

σ^b_{eq} satisfies equation (1.7b) as equality, and the first-best optimality condition (1.3) reduces to condition (1.8b). Hence, when condition (1.8b) holds as equality, the seller's equilibrium investment σ^b_{eq} satisfies the social optimality condition (1.3), and σ^b_{eq} is the socially optimal ex ante investment. When the left-hand side of condition (1.8b) exceeds its right-hand side, the marginal benefit side of the social optimality condition (1.3) exceeds the marginal cost side of condition (1.3), which implies an ex ante under-investment. Similarly, when the left-hand side of condition (1.8b) is less than its right-hand side, the marginal benefit side of equation (1.3) is less than the marginal cost side of equation (1.3), which implies an ex ante over–investment.

The question arises if the seller's equilibrium investment σ^b_{eq} can satisfy the social optimality condition (1.8b) for some choice of parameter values. To be specific, assume that the outside option is uniformly distributed, $\underline{V} \sim U[0, V]$. Following this assumption, we can replace $\lambda_1(\sigma^b_{sb}) + \lambda_2(\sigma^b_{eq}) = F_{\underline{V}}(0 < \underline{V} < \frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2}) = \frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{(1-\delta_2+\delta_1-\delta_1^2)V}$ and $\lambda_1'(\sigma^b_{sb}) = F'_{\underline{V}}(\underline{V} < \frac{\delta_1 w(\sigma)}{1+\delta_1}) = f_{\underline{V}}(\frac{\delta_1 w(\sigma)}{1+\delta_1})\frac{\delta_1 w'(\sigma)}{1+\delta_1} = \frac{\delta_1 w'(\sigma)}{(1+\delta_1)V}$ in equation (1.8b) and obtain equation (1.9b).

$$(1.9b)\ \frac{w(\sigma^b_{eq})}{w'(\sigma^b_{eq})} = \frac{\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{(1-\delta_2+\delta_1-\delta_1^2)V}}{\frac{\delta_1 w'(\sigma^b_{eq})}{(1+\delta_1)V}} \frac{\delta_1(1-\delta_1)(1+\delta_1)}{\delta_2-\delta_1}$$

Equation (1.9b) is satisfied if and only if the following condition holds:

$$(1.10b)\ 1 = \frac{\delta_1-\delta_1^2}{1-\delta_2+\delta_1-\delta_1^2} \frac{(1-\delta_1)(1+\delta_1)^2}{\delta_2-\delta_1}$$

Equation (1.10b) represents the relationship between depreciation rates δ_2 and δ_1 leading to the socially optimal equilibrium investment and it holds for any functions $w(.)$ and $c(.)$. Hence, the second-order concavity conditions A1.3.1 and A1.3.2 should be satisfied. If we solve the quadratic equation (1.10b) for δ_2, we obtain two solutions,

$$(1.11b)\ \delta_2^*(\delta_1)^{\pm} = \frac{1-\delta_1^2+2\delta_1}{2} \pm \sqrt{\frac{(1-\delta_1^2+2\delta_1)^2}{4} - \delta_1[(1-\delta_1)^2(1+\delta_1)^2 - \delta_1^2 + \delta_1 + 1]}.$$

The solutions converge for any $\delta_1 \in (0.25, 1)$. Next, we need to determine whether the socially optimal depreciation rates $(\delta_1, \delta_2^*(\delta_1)^{\pm})$ satisfy the restrictions imposed by the ex ante cutoff values, or whether $\delta_1(\delta_2^*) \in (\underline{\delta}_1^{EV}(\delta_2^*), \overline{\delta}_1^{EV}(\delta_2^*))$ and $\delta_2^*(\delta_1)^{\pm} > \underline{\delta}_2^{EV}(\delta_1, \delta_2^*(\delta_1)^{\pm})$.

Given the uniform ditribution of $\underline{V}$, we have $E[\underline{V}|\underline{V} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})] = \frac{1}{2}[\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2} + \frac{\delta_1 w(\sigma)}{1+\delta_1}] = \frac{1}{2}[\frac{\delta_1(2-\delta_2-\delta_1^2)w(\sigma)}{1-\delta_2+2\delta_1-\delta_1\delta_2-\delta_1^3}]$. Hence, $\underline{\delta}_1^{EV}(\delta_1,\delta_2), \overline{\delta}_1^{EV}(\delta_1,\delta_2) = \frac{1}{2} \mp \sqrt{\frac{1}{4} - \frac{\delta_1(1-\delta_2)(2-\delta_2-\delta_1^2)}{2-2\delta_2+2\delta_1-\delta_1\delta_2-\delta_1^3}}$ and $\underline{\delta}_2^{EV}(\delta_1,\delta_2) = 1+\delta_1-\delta_1^2 - \frac{2(1-\delta_1)(1-\delta_2+2\delta_1-\delta_1\delta_2-\delta_1^3)}{2-\delta_2-\delta_1^2}$.

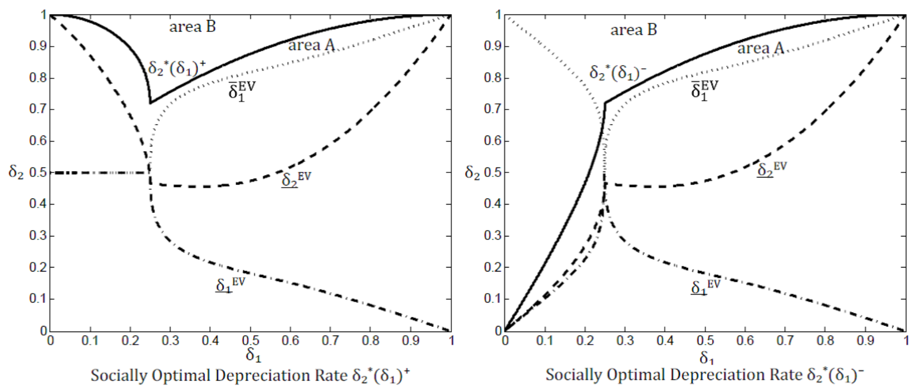

Figure 1.4. Socially Optimal Depreciation Rates when V˜$U[0, V]$

To see whether the restrictions on $(\delta_1, \delta_2^*(\delta_1)^{\pm})$ are satisfied we plot each socially optimal depreciation rate $\delta_2^*(\delta_1)$ along with the appropriate ex ante cutoff values $\underline{\delta}_1^{EV}(\delta_1, \delta_2^*(\delta_1))$, $\overline{\delta}_1^{EV}(\delta_1, \delta_2^*(\delta_1))$ and $\underline{\delta}_2^{EV}(\delta_1, \delta_2^*(\delta_1))$ in Figure 1.4. The graphs show that both plots of the socially optimal depreciation rates lie above their respective ex ante cutoff values. Hence, for each value of $\delta_1 \in (0,1)$ there exists at least one socially optimal depreciation rates δ_2^* such that each equilibrium specific investment is first-best optimal.

Figure 1.4 further shows that it is possible to have over-investment from the socially optimal point of view, if the pair of depreciation rates (δ_1, δ_2), lies in area B or above the graph of the socially optimal rate $\delta_2^*(\delta_1)$. Similarly, we can have under-investment from the socially optimally point of view, if the pair of depreciation rates (δ_1, δ_2), lies in area A or below the graph of the socially optimal rate $\delta_2^*(\delta_1)$. A pair of depreciation rates (δ_1, δ_2)

lying outside areas A or B results in the negotiation outcomes (a), which always result in under-investment.

An important conclusion from this example is that the resolution of the hold-up problem, when the outside option is uniformly distributed, does not depend on specific functional forms of the cost function $c(.)$ or the trade surplus function $w(.)$ as long as the second-order concavity conditions are satisfied. Further, the resolution of the hold-up problem critically depends on the relationship between depreciation rates δ_1 and δ_2. This is not surplising if we recall that the driving force behind different negotiation outcomes in the bargaining game is exactly the inter-relationship between depreciation rates on the outside option produced with generic investment and the tradable good produced with specific investment.

1.6 Discussion

The resolution of the hold-up problem in this chapter does not depend on timing, duration or a contractual form of interaction between trading partners. The key elements underlying the optimality of specific investment is the sufficient increase in the probability of the negotiation outcome λ_1, a sufficiently low depreciation rate on the buyer's outside option produced with generic investment ($\underline{\delta}_2^{EV} < \delta_2$) and a sufficiently high depreciation rate on the tradable good produced with specific investment ($\delta_1 < \overline{\delta}_1^{EV}$).

In contrast to other approaches to the hold-up problem, we find that the resolution of the hold-up problem does not depend on the seller's incentive to invest. As we show in Proposition 1.3, neither the improvement in the seller's ex ante negotiation position nor the presence of some specific depreciation rates are sufficient to induce the seller to increase specific investment ex ante. The resolution of the hold-up problem in our model is driven purely by the buyer's side. In particular, in the example with the uniformly distributed buyer's outside option we show that with an appropriate choice of depreciation rates we can have social optimality for any choice of the seller's specific investment, and that the social optimality of the equilibrium investment does not depend on functional forms of cost and payoff functions.

The hold-up problem is resolved because specific investment imposes a cost on the buyer by lowering the likelihood of the buyer's acquisition of an additional payoff due to appropriate depreciation rates. In particular, since specific investment raises the trade surplus in the relationship without changing the distribution of the parties' outside options, there is a fall in the likelihood that the buyer can extract the extra payoff by credibly threatening to reject the seller's offer. If the buyer's expected marginal cost from this fall in the likelihood is exactly matched by the buyer's expected marginal benefit from specific investment, then the seller's choice of specific investment is ex ante socially optimal.

Both key assumptions underlying the resolution of the hold-up problem are natural consequences of the specificity of investment. The assumption that the tradable good good with specific investment depreciates more than the outside option with generic investment reflects the fact that specific investment has no outside value, implying that there is no external market for goods produced with specific investment. Since the scrap value of the good with specific investment is virtually zero, we should expect a good with specific investment to depreciate at a higher rate than the good with generic investment.

The second assumption is the improvement in the seller's ex ante negotiation position from specific investment (the positive impact of investment on λ_1), and this assumption arises from the independence of specific investment and the distribution function of the buyer's outside option. Since specific investment raises the trade surplus in the relationship and does not affect the distribution function of the outside option, the likelihood that the buyer can extract a larger payoff by credibly using her threat of rejecting the seller's offer becomes smaller. Hence, with probability λ_1' the buyer loses the additional payoff above her outside option, which the buyer would have otherwise obtained should the seller not have undertaken specific investment.

The possible extensions to the model include the introduction of dynamics in investment decision. Since the first-best equilibrium investment in our static model is a subgame perfect Nash equilibrium, it should be first-best optimal in a repeated game as well. Moreover, if we invoke folk theorems we can achieve the first-best optimality of the seller's specific investment in a repeated game even if the seller's specific investment is not socially efficient in a static game. Thus, all our results should extend to models where the seller repeatedly undertakes specific investment. Another possible extension is the introduction of incomplete information about depreciation rates and the parties' payoff functions. We are not aware of any studies of incomplete information bargaining games with outside options, and an extension of the model in this direction may give more general results.

In general, the improvement in the seller's ex ante negotiation position depends on the ability of outsiders to make appealing counter-offers to the buyer after the seller undertakes investment. If outsiders can offer the buyer competitive counter-offers on par with the increased trade surplus in the relationship, the buyer can still maintain her prior negotiation position even after the seller undertakes investment. This ability depends on many factors aside from the degree of specificity of investment. The general competitiveness of the market, the ability of outsiders to replicate the specific investment of the seller, legal restrictions (patents and copyright laws), and the institutional environment also determine the degree to which the seller's ex ante negotiation position becomes stronger.

Another important implication of the model in this chapter is that the seller's equilibrium investment may lead to over-investment from the socially optimal point of view. There is anecdotal evidence of socially inefficient specific investment. For example, the instances of production of incompatible computer or electronic hardware by competing producers may serve as an example of over-investment. By engaging in trade of a generally incompatible device, a buyer faces costs in terms of a foregone or reduced ability to effectively use the threat of termination in the negotiation process, because termination would entail the extremely high expenses of a complete overhaul of the generally incompatible equipment.

The assumption that specific investment has no outside value is equivalent to the assumption that a tradable good with specific investment has no or few outside rivals. If by making specific investment the investor cannot hope to reduce the number of competing rivaling goods on par with the good produced by the seller, then the hold-up problem is likely to persist and the investor may under-invest ex ante. Based on this finding we propose three testable hypotheses. Firstly, we argue that we should observe more severe under-investment in relationship-specific assets in markets with highly competitive environments and lower or no under-investment in markets with less competitive environments. The next empirical hypothesis is related to the duration of relationships with specific investment. We claim that because specific investment lowers the probability of termination of the relationship, the duration of relationships with specific investment should be longer. There are already many studies which largely support this hypothesis[6]. Lastly, we predict an increase in the investor's bargaining power from specific investment in markets with less competitive environments. This hypothesis is at odds with the traditional conclusion, which suggests that in general the investor's bargaining power should become weaker due to investment specificity. To distinguish between our hypothesis and the traditional conclusion we would like to emphasize the importance of the competitive environment in formulating our hypothesis.

1.7 Concluding Remarks

In this chapter we demonstrate that the hold-up problem does not necessarily plague all relationships with specific investment. The hold-up problem can be absent depending on the degree to which specific investment adjusts the parties' ex ante negotiation positions

[6]For example, Joskow (1985) finds that specificity of investment has a positive effect on the contractual duration in his study of contractual relationships between coal suppliers and electricity generating plants.

and the relationship between depreciation rates on the tradable good with specific investment and the outside option produced with generic investment. We further demonstrate that it is possible to completely resolve the hold-up problem with an appropriate choice of parameters.

In addition, we characterize a set of negotiation equilibria in the Rubinstein's alternating-offers bargaining game with an outside option. We extend the results of Rubinstein (1982) and Binmore et. al. (1989) by introducing differentiated depreciation rates on the outside option and the tradable good, and show how the new negotiation outcomes lead to the first-best optimality of specific investment.

Chapter 2

Auctions as Appraisal Mechanisms: Seller Behavior in eBay Auctions

In this chapter we study seller behavior in eBay auctions for high-valued goods where external appraisal values are difficult or costly to obtain. We find that many sellers use eBay auctions not only to sell such items but also to gather information about items' valuations from observed bids in eBay auctions. Sellers use the gathered information to update their beliefs about distribution functions of items' valuations. To gather information about items' valuations from realized bids without necessarily selling the items, sellers in eBay auctions employ high secret reserve prices, because high secret reserve prices prevent auction sales and at the same time do not restrict entry of bidders into auctions. To signal the quality of their items many sellers use public reserve prices together with secret reserve prices. The proposed behavior of sellers on eBay explains several previously unanswered stylized facts in the auctions literature: the widespread use of secret reserve prices by sellers, the presence of repeated auctions, and the low rate of success of eBay auctions for some item categories.

Like other market transaction mechanisms, an auction transfers an item from a seller to a buyer at some specified price. However, unlike other market transaction mechanisms, an auction, and in particular an ascending-price auction, also acts as an information gathering mechanism by inducing competition among buyers and by forcing these buyers to reveal their valuations of the auctioned item. If a seller has no clear idea about the value of her item, is sufficiently patient, and wants to sell at the highest possible price, she clearly gains by running at least one ascending-price auction without selling the item and learning the distribution of realized bids before engaging in an actual sale. The ability to observe realized bids without selling an item can help this uninformed and patient seller to not only discover the demand for her item but also to choose an optimal selling price. In particular, the uninformed and patient seller can run an unsuccessful ascending-price auction, learn the highest bid, and then offer her item in a posted price sale at the highest bid from the unsuccessful auction stage. This way the seller's expected revenue should equal the expected highest valuation minus the cost of an unsuccessful auction. If the seller's cost of an unsuccessful auction is sufficiently small, then the expected payoff from implementing this two-stage sale clearly exceeds the expected second highest valuation - an expected payoff from running a single-stage ascending-price auction.

The question is whether an uninformed and patient seller who wants to learn the distribution of an item's valuation can convince buyers to enter and bid truthfully in the auction stage(s). Electronic selling platforms such as eBay offer an option suitable for this purpose. The rules of an eBay auction allow a seller to set a secret reserve price hidden from bidders. Bidders do not observe a secret reserve price itself, but they can see if a secret reserve price is set until some bid exceeds it. By setting a high secret reserve price, a seller can run one or multiple unsuccessful ascending-price auctions, observe the truthfully revealed valuations, and use this knowledge either in a different selling format on eBay or delist her item from eBay and use the acquired knowledge in a sale elsewhere[1]. In this chapter we provide empirical evidence that secret reserve prices could be used exactly for this purpose: as an effective instrument, which allows uninformed and patient sellers to use eBay auctions not only as a selling mechanism but also as an affordable alternative to an otherwise expensive or unavailable appraisal mechanism.

This chapter fits into the literature on the empirical estimation of auctions data. However, unlike most studies, we concentrate on the behavior of sellers rather than on the

[1]Specifically, an uninformed and patient seller can use a posted price sale format available on eBay after running an unsuccessful eBay auction and learning parameters of the distribution function of bidders' valuations.

behavior of bidders[2]. In particular, we study how and why sellers set public and secret reserve prices in eBay auctions.

Many studies test public and secret reserve prices in the context of eBay auctions. However, most of them test secret and public reserve prices only in static auctions and ignore the fact that many unsold items are relisted. In general, existing empirical studies find that public reserve prices raise sale prices, while secret reserve prices have either a negative or an insignificantly positive effect on sale prices. For example, Bajari and Hortacsu (2003) study public and secret reserve prices in static eBay auctions for collectible US coins and find a small positive effect of secret reserve prices on sale prices. Katkar and Reiley (2006) employ field experiments and use static eBay auctions to sell Pokemon cards with public and secret reserve prices. They find that the presence of a secret reserve price lowers revenue by 10% and the likelihood of sale by 34%. Lucking-Reiley et. al. (2007) use data from static eBay auctions of collectible pennies and find that, conditional on sale, public reserve prices have a small positive (around 1%) but statistically insignificant effect on final sale prices while secret reserve prices have a sizeable positive (about 15%) and statistically significant effect on final sale prices[3]. Carare (2012) studies revenue effects of only public reserve prices in dynamic auctions for computer processors and finds that optimal public reserve prices in dynamic auctions have a significant positive effect on seller revenue[4].

The theoretical discussion of the functions and determinants of public reserve prices depends on whether bidders' valuations are independently distributed or have a commonly distributed component. If bidders' valuations are independently distributed, then, according to Myerson (1981) and Riley and Samuelson (1981), sellers use a public reserve price to screen out bidders with valuations below some threshold level. If bidders' valu-

[2]In a recent working paper Einav et al. (2012) also address seller behavior in online markets. They argue that sellers use online platforms such as eBay to experiment with auction parameters and with sale formats.

[3]Lucking-Reiley et. al (2007) study sale effects of public and secret reserve prices only in successful auctions without relistings.

[4]Carare's derivation of an optimal public reserve price is different from the standard Myerson's optimal public reserve price, because Carare's derivation accounts for the possibility of a future relisting.

ations have a commonly distributed component, then as Milgrom and Weber (1982) and Cai, Riley and Ye (2007) show, a public reserve price serves as a credible signal of the quality of an item on sale.

While the functions and determinants of public reserve prices have been thoroughly researched and understood in auction theory, theoretical studies of auctions are silent on the function and determinants of secret reserve prices. We are aware of only two theoretical studies of secret reserve prices, and both these studies look at secret reserve prices as possible alternatives to public reserve prices. Vincent (1995) shows that if bidders' valuations are commonly distributed, then secret reserve prices can generate more revenue than public reserve prices by encouraging entry. Rosenkranz and Schmitz (2007) use prospect theory and find that if bidders' valuations are independently distributed, a public reserve price enters bidders' utility functions, and if bidders' outside options exceed the public reserve price, then a secret reserve price can outperform a public reserve price. However, no empirical analysis of either of these theories is present in the literature.

We can draw several conclusions from the existing theoretical and empirical studies of seller behavior in auctions. First of all, there is a gap in the theoretical literature related to the use and the function of a secret reserve price. The available theoretical studies attribute functions of a public reserve price to a secret reserve price and analyze the effectiveness of a secret reserve price as an alternative to a public reserve price. We are able to address this gap in the literature by relaxing the theoretical assumption that sellers have a perfect knowledge of the distribution function of bidders' valuations and by showing that an uninformed patient seller can use a secret reserve price to obtain information about the distribution function of valuations by running an unsuccessful ascending-price auction. Hence, we argue that a secret reserve price has a completely different purpose than a public reserve price. While a public reserve price is used to screen out low-valuation bidders and to signal the quality of an item on sale, a secret reserve price is primarily used by an uninformed and patient seller to collect information about the distribution function of valuations prior to engaging in an actual sale.

The existing empirical studies of seller behavior in auctions are inconclusive about the impact of secret reserve prices on seller revenue and sale prices. We test revenue effects of secret reserve prices under the assumption that secret reserve prices are used not as screening or signaling devices but as a mechanism through which uninformed and patient sellers are able to gather information about items' valuations. We find that under our specification, secret reserve prices have a strong positive effect on sale prices. We run a simulation exercise using the actual data from sales of used tractors in eBay auctions, where we test the price effect of imposing a secret reserve price in repeated auctions under different belief-updating rules. We find that as long as the cost of relisting is not too high and an uninformed seller updates a secret reserve price given the information from previous unsuccessful auctions, the sale price in an auction with an appropriately set secret reserve price exceeds the sale price in an auction without a secret reserve price[5].

Our final contribution is related to determinants of public reserve prices, secret reserve prices and buy-it-now prices on eBay. We are not aware of any empirical studies with structural tests of determinants of secret and public reserve prices using eBay data. We find that sellers of used tractors on eBay use public reserve prices both to screen out low valuation bidders and to signal the quality of listed tractors. Hence, many empirical studies, which a priori assume that public reserve prices are used exclusively for screening, overlook an important component defining seller behavior. We also test for various determinants of secret reserve prices and find that the size of a secret reserve price is determined by the highest bid from previously run unsuccessful second-price auctions and by average highest bids observed in auctions for similar items. In addition, we test for determinants of buy-it-now (BIN) prices and find that buy-it-now prices are almost exclusively determined by highest bids from previous unsuccessful auctions. To test for determinants of public reserve

[5]In the simulation exercise we account for the possibility that an item may go unsold if the secret reserve price is not met after multiple relistings. We assume that the sale price is zero if the item does not sell after the fifth relisting.

prices, secret reserve prices, and BINs, we employ a completely non-parametric approach to minimize the number of identification assumptions.

In the next section we discuss the data. In section 2.2 we define the model of a seller and a bidder behavior. In section 2.3 we discuss the estimation strategy. In section 2.4 we discuss determinants of public reserve prices, secret reserve prices, and BINs. In section 2.5 we test determinants of sale prices.

2.1 Data

To study the behavior of sellers on eBay we use the data on used tractors sold in eBay auctions between 11/17/04 and 5/30/07. Our data consists of two components: (1) a complete account of all auction sales on eBay between 11/17/04 and 5/30/07 without bidder characteristics, and (2) bidder characteristics and actual bids for a portion of auctions held on eBay between 11/17/04 and 5/30/07. Since we don't have realized bids and bidder characteristics for a majority of auctions held on eBay, for the analysis of a seller behavior we construct a smaller sample of auctions with complete information on bids and bidder characteristics[6]. To give an idea about the market for tractors on eBay, we first present information about all auction sales of tractors held on eBay between 11/17/04 and 5/30/07. However, when we discuss tractor characteristics, we present information only about those tractors, for which we have complete information on bids and bidder characteristics[7].

A cursory look at outcomes of sales of tractors in eBay auctions suggests that many auctions are not successful and a vast number of tractors are not sold on eBay. This is a mystery, especially, if we recall that auctions have an advantage over other sale formats

[6]There is a sizeable share of auctions with missing information about tractor characteristics such as horse power and year of production. We use tractor model numbers to recover these parameters from outside sources. The main source of the outside data is www.tractor-data.com.

[7]All the auctions with missing bidder characteristics are auctions taking place in 2006 and 2007. We find a bias in the sub-sample of auctions with known bidder characteristics: for example, the average tractor horse power in the sub-sample of auctions with known bidder chracteristics is less than the average horse power in the full sample by 1.43 HP. This difference is statistically significant at 1% level with t-statistic of 15.421.

(such as posted price sales) in terms of a higher probability of success. There is a reasonable suspicion that many of these unsold tractors are relisted over and over again. We are able to match tractors from multiple listings and construct a dataset of tractors with multiple relistings[8]. We present the data on all repeated and single auction sales of tractors held on eBay between 11/17/04 and 5/30/07 in Table 2.1.

Listing patterns	Number listed	Number sold	Percent sold
Single listing	23253	13251	56.99
2 listings	4031	1448	35.91
3 listings	1069	344	32.19
4 listings	404	117	28.96
5 listings	197	60	30.46
6 listings	102	30	29.41
7 listings	61	20	32.79
8 or more listings	111	34	30.63
Total number of unique tractors	29228	15304	52.36

Notes: Single listing tractors might have been listed multiple times prior to 11/17/04, when the earliest observations in our data first appear.

Table 2.1. Auction sales of tractors on eBay between 11/17/04 and 5/30/07

We can notice two regularities by looking at Table 2.1. First, many tractors are delisted from eBay auctions after being unsuccessfully auctioned[9]. In fact, only about a half of all

[8]When deciding whether it is the same tractor with multiple listings or different tractors, we use information on seller id's, whether the prior sale was successful, tractor characteristics, and engine hours.

[9]By "delisted" we mean unsuccessfully auctioned tractors that are not sold in auctions again. It is possible that some unsuccessfully auctioned tractors are relisted on eBay under a different sale format (such as, for example, an eBay posted price sale). In this dataset we do not have information on eBay posted price sales.

tractors are successfully sold even if we account for relistings. The highest share of successfully sold tractors is found among tractors which were listed only once[10]. The second interesting regularity is that the share of sold tractors across multiple relistings is stable and on average does not diverge from a 30% success rate.

Next, we present the data on times between listings for tractors listed twice. Table 2.2 shows that a half of unsuccessfully sold tractors are relisted within 3 days after the end of an unsuccessful sale and almost 30% of unsuccessfully sold tractors are relisted within the same day. The maximum time between listings for tractors listed twice is 540 days, or about a year and a half. The average time between listings is 17.36 days.

Duration between relistings	Number of 2-stage auctions	Percent
less than 1 day	1201	29.79
1-3 days	811	20.12
4-7 days	559	13.87
8 days to 2 weeks	439	10.89
2 weeks to 1 month	525	13.02
from 1 month to 2 months	262	6.50
more than 2 months	234	5.81
Total	4031	100.00

Table 2.2. Times between listings for tractors listed twice

Sellers on eBay can choose the duration of their auctions. In our data the maximum duration of auctions is 10 days and the minimum duration of auctions is less than 1 day. More precisely, 30% of auctions last less than 7 days, 53% of all auctions last exactly 7 days and about 18% of auctions last 10 days.

[10]Note that tractors with 1 listing might have been listed multiple times prior to 11/17/04, when the earliest observations in our data first appear.

We present the data on various auction parameters used by sellers of tractors on eBay in Table 2.3. Sellers on eBay have a choice to set three parameters before running an auction: a public reserve price (PRP), a secret reserve price (SRP), and a buy-it-now price (BIN)[11]. A public reserve price defines the minimum bid from which participating bidders must start. A secret reserve price defines the unobservable reserve price, which the highest bid must exceed to successfully end an auction. If the highest bid is below a secret reserve price by the end of an auction, the item goes unsold and the auction ends unsuccessfully. Although a secret reserve price is not observable, bidders can see an indicator "Reserve Not Met" if there is an unmet secret reserve price in an eBay auction listing. The secret reserve price indicator disappears, once a bid exceeds it. A seller can adjust her secret reserve price at any time before the last 12 hours of an active auction. In the last 12 hours of an auction an active secret reserve price cannot be changed or deactivated.

Auction parameter	Observability	eBay fee for tractors
Public reserve price	Fully observable	\$5
Secret reserve price	Unobservable price, observable indicator	\$5
Buy-it-now price	Observable prior to auction, unobservable later	if listing $\leq$ 50 items/month, free, otherwise \$0.25/listing

Table 2.3. Parameters in eBay auctions

A BIN is a publicly observable temporary sale price at which any bidder can purchase an item before the start of an auction. Any indicator of a BIN or its presence disappears

[11]On eBay a public reserve price is called a starting price, while a secret reserve price is called a reserve price. To avoid confusion, in the rest of the paper we adopt the standard terminology from auctions literature rather than the terms used on eBay.

in an auction without a secret reserve price, once any bidder makes a bid. If a BIN is used together with a secret reserve price, then the BIN indicator is observable until a secret reserve price is met. We present a description of these three parameters and fees for using them in auctions for tractors[12].

In the next two tables we present the data on the use of different auction parameters across sales with one and two listings. We can see that at least a third of sellers use a public and a secret reserve price at the same time. Further, more than a half of sellers of tractors with two listings and 46% of sellers of tractors with a single listing employ secret reserve prices alone or in combination with a public reserve price and a BIN. We can also see that the use of secret reserve prices is more than 6% higher in the first listing than in the second listing for tractors listed twice and that the use of a public reserve price and a secret reserve price in the first listing (all unsuccessful auctions) among tractors listed twice is much higher than the use of a public and a secret reserve price for tractors listed only once.

[12]eBay fees for using public/secret reserve prices and BINs depend on an item category.

Features used	Number	Percent
Auctions with a public reserve price (PRP) (>$100)	15573	66.97
Auctions with a secret reserve price (SRP)	10705	46.04
Auctions with a buy-it-now price (BIN)	5208	22.39
Auctions with PRP and SRP	7544	32.44
Auctions with PRP, SRP and BIN	1925	8.28
Auctions with SRP and BIN	2861	12.30
Auctions without any features	4135	17.78
Total number of tractors listed once	23253	100.00

Table 2.4. The use of different auction features for tractors listed only once

Features used	First listing		Second listing	
	Number	Percent	Number	Percent
Auctions with a public reserve price (PRP)(>$100)	3763	93.35	3301	81.89
Auctions with a secret reserve price (SRP)	2314	57.41	2064	51.20
Auctions with a buy-it-now price (BIN)	1250	31.01	1375	34.11
Auctions with PRP and SRP	2094	51.95	1506	37.36
Auctions with PRP, SRP and BIN	643	15.95	533	13.22
Auctions with SRP and BIN	724	17.96	751	18.63
Auctions without any features	45	1.12	136	3.37
Total number of tractors listed twice	4031	100.00	4031	100.00

Table 2.5. The use of different auction features for tractors listed twice

In the next table and two figures we present the data on tractor characteristics, seller ratings, and the number of bidders in those auctions, for which we have complete information about all realized bids and all bidders. Table 2.6 shows that the average number of bidders in auctions with complete bidder information is about 8.5 with a minimum of 1 and a maximum of 32. Since we don't have complete information on bids for all auctions, we cannot identify auctions with zero entry from auctions with missing bidder characteristics. This is why we do not have auctions with zero entry in our sample and the minimum number of bidders in our sample is 1.

	Number of observations	Mean	Standard Deviation	Min	Max
Number of bidders	12429	8.469	5.282	1	32
Seller feedback score	12429	201.652	937.296	0	80282
Tractor age	12429	22.212	16.749	0	80
Tractor horse power	12429	42.076	29.446	10	150

Table 2.6. Tractor characteristics in auctions with complete bidder information

Table 2.6 shows that the average age of tractors in our data is more than 22 years and 68% of tractors range in age from 5.5 years to 38.9 years. The distribution of tractors by age is given in Figure 2.1. By looking at the figure we can see that there are two clearly identified modes. The first mode is located at the age of 1 and the second mode is located at the age of 25. We think that the bimodal distribution of tractor ages can be explained by practices of manufacturers of tractors. Manufacturers of tractors usually update their production lines from 4 to 10 years. As a result, sellers of used tractors on eBay have a stronger incentive to sell their tractors once new generations of tractors become available. The mode at the age of 1 clearly captures sales of the latest generations of tractors, while the mode at the age of 25 captures sales of tractors of previous 2 or 3 generations, when the difference between tractor characteristics from the current generation and the past generations becomes more critical.

Figure 2.1. Age of Tractors

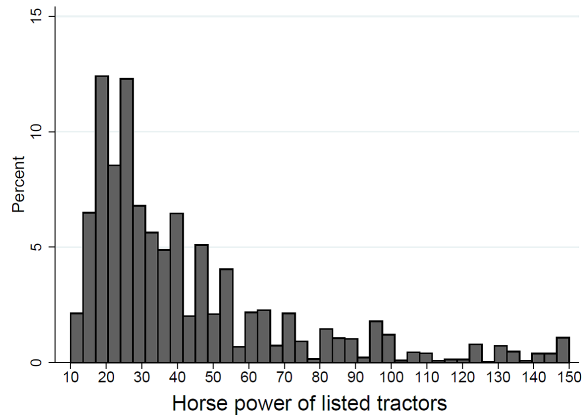

Figure 2.2. Tractor HP

We present the distribution of tractors by horse power in Figure 2.2. In general, tractors can be separated into three categories by horse power: lawn tractors, small or compact utility tractors and utility tractors. Lawn tractors rarely exceed 20 HP in power and are used for mowing lawns or carrying relatively light loads. Small or compact utility tractors range in power from 20 HP to 40 HP and are the most popularly used tractors for everyday farming needs. Utility tractors range in power from 40 HP to 300 HP and are used for cropping, construction, and other heavy duty tasks. The distribution of tractors by horse power in Figure 2.2 shows that the most frequently listed tractors are either lawn tractors with HP below 20 or compact utility tractors with HP between 20 and 40.

2.2 Model and Predictions

In this section we discuss the model of bidder and seller behavior on eBay. The bidder behavior in eBay auctions has been extensively discussed (see Hasker and Sickles (2010) for the latest survey), and it can be described as "proxy bidding." To participate in proxy bidding in an eBay auction, a bidder has to specify the maximum willingness to pay. A computer then raises bids in increments from the minimum starting price set by the seller and up to the value indicating the maximum willingness to pay specified by the bidder. A winning bidder is notified by E-mail about the winning price after the end of an auction. Bidders can raise values of their willingness to pay or to make bids personally at any time during an auction.

Given the rules of proxy bidding, the key feature of an eBay auction is that bidders incur minimum participation costs. Since bidding on eBay does not require bidders' physical participation and is free, a typical bidder does not have to face any participation costs aside from the search costs.

Since bids increase and a winning bidder pays an increment above the second highest bid in an eBay auction, we use a theoretical model of a static ascending second-price auction to model bidder behavior[13]. We further assume that bidders' valuations have a finite common support $[0, v]$ and are identically and independently distributed. The independence assumption is used to capture the fact that we study a market of used goods with an unlikely possibility of a post-auction resale[14]. We further allow that bidders engage in "sniping," or that bidders place their bids in the last minutes of an auction closing time[15]. According to Ockenfels and Roth (2006), the main consequence of sniping is that some bidders are not able to place their highest bids before an auction closing time. As a result, the realized bids of non-winning bidders may not reflect the bidders' valuations of an item on sale.

According to bidding strategies in an ascending second-price auction with sniping, the winning bidder in an eBay auction for used tractors bids the second-highest valuation or

[13]By the static ascending second-price auction we mean a standard second-price auction without resale or relisting and satisfying the assumptions of Myerson (1981).

[14]Unfortunately, it is impossible to identify if tractors are bought for resale, because we do not observe tractors' serial numbers and we do not have data on the fate of purchased or unsold tractors outside eBay.

[15]We find that 36.23% of auctions in our sample have at least one bid placed in the last minute of an auction closing time.

some value below the second-highest valuation[16]. All other bidders bid their true valuations or some values below them. As a result, if an eBay auction is unsuccessful and there is no winner, the seller and the bidders observe true valuations of all participating bidders or some values below them, and if an auction is successful and the item is sold, the seller and the bidders observe true valuations or some values below them of all participating bidders but the winner. Hence, we state our first two assumptions about bidder behavior in eBay auctions.

Assumption 2.1. Bidders' cost of participation is zero.

Assumption 2.2. Bidders' valuations are identically and independently distributed on a finite support $[0, v]$.

It is harder to formalize seller behavior on eBay, because in the theoretical auction literature seller behavior is analyzed under the same informational assumptions as bidder behavior. In particular, it is usually assumed that seller and bidder valuations are realizations of the same distribution function, and that the seller and bidders know this distribution function but do not know each others' realized valuations. Under the assumption that the seller knows the distribution function of bidders' valuations, the sole role of a profit-maximizing seller is to calculate and set an optimal public reserve price before the start of an auction.

As we can see the actual eBay auction format gives a much richer set of actions to a seller. In addition to choosing a sale format, which is beyond the scope of discussion of this chapter, in a standard eBay auctions a seller can set a public reserve price, a secret reserve price and a buy-it-now (BIN) price.

The function of a public reserve price depends on whether bidders' valuations are independently distributed or have a commonly distributed component. If bidders' valuations

[16]Due to sniping, there is a positive probability that the second-highest bidder is not able to bid her true valuation and the winning bid is below the second-highest valuation.

are independently distributed, then according to Myerson (1981) and Riley and Samuelson (1981), a public reserve price screens out low valuation bidders. If bidders' valuations have a commonly distributed component, then Cai, Riley, and Ye (2007) show that a public reserve price is a credible signal of the value of the item on sale.

The second parameter available to a seller is a secret reserve price. An important difference between a secret reserve price and a public reserve price is that a secret reserve price does not restrict bidders' entry, and bidders with valuations below a secret reserve price can freely participate in an auction.

The last parameter available to a seller on eBay is a buy-it-now (BIN) price. A seller can set a price at which the item on sale can be sold immediately before the beginning of an auction. Once the first bid greater than the secret reserve price (when applicable) is made, the BIN and any indicator of its presence disappear and no bidder can observe the BIN later.

Before we continue discussing seller choice of reserve prices and a BIN, we need to specify how bidder behavior is affected by these parameters. Before we proceed, we assume that bidders and the seller are risk-neutral.

Assumption 2.3. Bidders and the seller are risk-neutral.

The impact of a public reserve price on bidder behavior has been extensively discussed in the literature, and under assumption 2.2 it amounts to limiting the participation of bidders with valuations below some threshold level. If assumption 2.2 is violated and bidders' valuations have a commonly distributed component, then a public reserve price signals the quality of an item on sale and raises bids of participating bidders. Nevertheless, whether assumption 2.2 is violated or not, when a public reserve price is present, no bidder with a valuation below the public reserve price should enter an auction.

Mathews (2003) has discussed the optimality of temporary BINs used on eBay and the effect of such BINs on bidder behavior[17]. He shows that in a particular case when bidders and a seller are risk-neutral, there is an equilibrium in which a profit-maximizing seller sets a BIN equal to the upper bound of the support of the distribution function of bidders' valuations, while bidders bid without taking the BIN option. When the seller is risk-averse, in equilibrium the seller sets the BIN below the upper bound of the support of valuations, while bidders take the BIN option with a positive probability. Using Mathews' result, we conclude that under risk-neutrality, the presence of a BIN in an eBay auction should not affect bidders' participation rates and bidding strategies.

However, whenever a BIN is used together with a very high secret reserve price, a temporary BIN effectively turns into a permanent BIN, since by eBay rules a BIN is active as long as the secret reserve price is not met. In this case, Hidvegi et. al. (2006) show that a bidder's equilibrium strategy depends on whether her valuation is above or below an active BIN. If the bidder's valuation is below an active BIN, then the bidder should truthfully bid up to her valuation. If the bidder's valuation is above an active BIN, then the bidder should bid truthfully until winning an auction or until reaching some threshold value after which the bidder should take the BIN.

The theoretical discussion of a secret reserve price and its effect on bidder behavior has received only limited attention in the theoretical literature. The available studies view a secret reserve price as an alternative to a public reserve price, and assume that a secret reserve price plays the same screening role as a public reserve price, however, without limiting the entry. For example, Vincent (1995) argues that if bidders are risk-averse, then in a common-value environment a secret reserve price may have an advantage over a public reserve price by encouraging entry.

[17] The difference between a temporary BIN used in eBay auctions and a permanent BIN used in other selling formats is that a temporary BIN disappears once an auction starts, while a permanent BIN stays available during the course of the whole auction. For discussion of a permanent BIN, see, for example, Hidvegi et. al. (2006).

Based on assumption 2.1 that bidders' participation costs in a standard eBay auction are zero, we also argue that the presence of a secret reserve price should not affect bidders' entry as long as there is a non-zero probability of winning an auction. The next question is whether participating bidders have a different bidding strategy in the presence of a secret reserve price. The answer to this question depends on whether participating bidders acquire any additional information by observing the presence of an active secret reserve price and whether bidders obtain any additional payoff in an auction with a secret reserve price.

We argue that the presence of a secret reserve price in a single-stage ascending second-price auction does not reveal any additional information to bidders nor does it give any additional payoff above the standard expected bidder payoff[18]. Hence, if bidders participate in a single-stage ascending second-price auction with a secret reserve price, we should not observe any changes in the bidder behavior from an equilibrium bidder strategy in a single-stage ascending second-price auction without a secret reserve price. However, if the same bidders participate in several auctions for the same item and with secret reserve prices in one or more listings, the informational structure of bidders and bidders' expected payoffs should change. In particular, if bidders participate in a repeated ascending second-price auction with a secret reserve price in each stage, (a) the bidders are able to update their beliefs about the secret reserve price at the end of each unsuccessful stage, and (b) the bidders should take into account an expected discounted payoff from participating in all auction stages under an evolving information structure. To simplify analysis, we assume that bidders do not participate in more than one stage of an eBay auction with multiple relistings.

Assumption 2.4. Bidders participate only in one stage of auctions with multiple relistings.

[18]By a single-stage ascending second-price auction we mean a standard static second-price auction without resale or relisting and satisfying the assumptions of Myerson (1981).

Given assumptions 2.1, 2.2, 2.3, and 2.4 we are ready to define bidder strategies in an eBay auction with one or more relistings in the presence of a public reserve price, a secret reserve price, and a BIN.

Proposition 2.1. *Let* $[0, v]$ *be a finite support of the distribution function of bidders' valuations,* $r \geq 0$ *be a public reserve price,* $s \geq 0$ *be a secret reserve price,* $b \geq 0$ *be a buy-it-now price, and N be a number of bidders. Then under assumptions 2.1, 2.2, 2.3, and 2.4, the following holds:*

a) Conditional on a bidder's valuation exceeding the public reserve price r*, it is a weakly dominant strategy for the bidder to enter an auction if the lower bound of the support of a secret reserve price is the public reserve price* r*.*

b) Equilibrium bidding strategies in an eBay auction with r,s,b, N bidders and sniping coincide with equilibrium bidding strategies in an eBay auction with r, without s, N+1 bidders, and sniping.

In the proof we concentrate on bidders' decision to participate only in the presence of a secret reserve price. For the discussion of bidders' strategies in the presence of a public reserve price, see Myerson (1981). For the discussion of bidders' strategies in the presence of a temporary BIN and a permanent BIN, see respectively Mathews (2003) and Hidwegi et. al. (2006).

Proof:

Under symmetry of distribution functions of bidders' valuations, let N be the number of bidders in an auction, $v_i \in [0, v]$ be a valuation of a representative bidder i, $b_i \in [0, v_i]$ be an equilibrium bid of bidder i, and assume that $v_i \geq r$. Since bidder i does not observe a secret reserve price, bidder i treats the secret reserve price s as a random variable with a support $[\underline{s}, \overline{s}]$. Under the assumption that the lower bound of the support is $\underline{s} = r$, costless bidding, and if bidder i follows an equilibrium bidding strategy in a second-price auction as in Myerson (1981), bidder i's expected payoff from entry is $(v_i - b_i)\text{Prob}(b_i \geq b_{j \neq i}, b_i \geq s, b_i \geq r) \geq 0$. If bidder i does not enter, then bidder i's payoff is zero. Hence, it is a weakly dominant strategy for bidder i to enter an auction if the lower bound of the support of the secret reserve price s is $\underline{s} = r$.

To see that part (b) is true, note that the formulation of the problem with a secret reserve price, a public reserve price and N bidders in Proposition 2.1 is identical to the formulation of a problem with a public reserve price and $N+1$ bidders, where the $(N+1)$st bidder's valuation is $s \in [\underline{s}, \bar{s}]$. Hence, bidder i's equilibrium bidding strategy in an ascending second-price auction with sniping, N bidders, a public reserve, and a secret reserve coincides with an equilibrium bidding strategy in an ascending second-price auction with sniping, $N+1$ bidders, a public reserve, and no secret reserve. ■

Given the equilibrium strategies of bidders in the presence of a secret reserve price, we can discuss the seller equilibrium behavior. We relax the standard assumption that the seller knows the distribution function of bidders' valuations and assume that the seller neither perfectly knows the distribution function of bidders' valuation nor knows its support.

Assumption 2.5. A seller does not have perfect knowledge of the distribution function of bidders' valuations and its support.

If the seller does not know the distribution function of bidders' valuations, then she cannot set an optimal public reserve price and an optimal BIN. Hence, it is natural to expect that the seller may want to gather some information about the distribution function of bidders' valuations, if the cost of such information is sufficiently low and the future discount factor is one. The natural place to gather such information is an ascending auction itself. An ascending-price auction is the best source of information about the distribution function of bidders' valuations as opposed to a descending-price auction or a posted price sale, because in an ascending-price auction unsuccessful bidders reveal the most accurate information about their valuations. The main drawback of an ascending auction is that the winning bidder may not reveal her valuation[19]. However, if the item does not sell and there is no winner, even in the presence of sniping a seller is able to gather the most accurate information about valuations of all participating bidders. Hence, if the cost of running an unsuccessful ascending-price auction is small and an uninformed seller cares only about learning the distribution function of bidders' valuations, then the seller has an incentive to run an unsuccessful ascending-price auction to learn the distribution function of bidders' valuations to form a belief about the distribution function of bidders' valuations.

The use of an ascending-price auction in this fashion relies on three key elements: the low cost of an unsuccessful auction in terms of a physical cost and foregone time, a sufficient entry in an auction, and the ability to run an auction without selling an item. All three elements are available to a seller on eBay. The cost of running an eBay auction varies across item categories. In our particular case, the cost of listing a tractor is $20, the cost of using a public or a secret reserve price is $5, and the cost of using a BIN never exceeds $0.25. Hence, the total cost of unsuccessfully auctioning a tractor on eBay with a secret reserve price at most amounts to $25.25[20].

The assumption that sellers who use auctions to learn distributions of valuations do not discount the future is based on the observation that we should observe a self-selection of patient sellers into such auctions, since rational sellers should expect a low probability of success in auctions with secret reserve prices. This claim is further reinforced by the empirical finding that about 63% of relists occur within 7 days after an unsuccessful listing. Given these considerations we are ready to state our next assumption.

Assumption 2.6. A seller has the future discount factor of 1 and faces negligible costs of running an unsuccessful ascending-price auction on eBay.

[19]It is a weakly dominant strategy to reveal valuation for the winner only in an ascending sealed-bid second-price (Vickrey) auction.

[20]The cost of an unsuccessful auction with a secret reserve price does not include a final value fee, which a seller has to pay if the item sells. For tractors sold in eBay auctions the final value fee is %1 of the sale price with a maximum of $250.

The second key element is that bidders enter an eBay auction with a secret reserve price and always bid according to equilibrium strategies of a single-stage ascending second-price auction with sniping. By part (a) of Proposition 2.1, it is a weakly dominant strategy to enter an auction with a secret reserve price, and by part (b) of Proposition 2.1, participating bidders use equilibrium bidding strategies of a single-stage ascending second-price auction with sniping.

Lastly, the ability of an uninformed seller to run an ascending-price auction without selling an item relies on the availability of a secret reserve price in eBay auctions. In particular, the ability to set an unmet secret reserve price allows an uninformed seller to run an ascending second-price auction without selling an item and without limiting entry. In the next proposition we state how a seller should set an equilibrium secret reserve price.

Proposition 2.2. *Let $[0, v]$ be an unknown support of the distribution function of bidders' valuations, $s \geq 0$ be a secret reserve price, and $b \geq 0$ be a BIN price. Then under assumptions 2.1, 2.2, 2.3, 2.4, 2.5, and 2.6, in equilibrium an uninformed seller, who does not discount the future, sets a secret reserve price equal to the maximum of the highest order statistic of the distribution function of bidders' valuations.*

Proof: First, assume that the seller knows the support of the distribution function $[0, v]$ but does not know the distribution function itself. Then, according to Mathews (2003), a risk-neutral seller should set a temporary BIN equal to v, while a risk-averse seller should set a temporary BIN to some value below v. Hidvegi et al. (2006) show that whether a seller and bidders are risk-neutral or not, a seller should set a permanent BIN above or equal to the expected highest valuation of a symmetric bidder.

Next, by Proposition 2.1, the most complete revelation of bidders' valuations in an ascending-price auction is achieved if the seller sets a secret reserve price s above or equal to v, $s \geq v$. Under the assumption that the distribution function of bidders' valuations is discrete, the seller obtains a strictly positive expected payoff by setting the secret reserve s equal to v. If the distribution function of bidders' valuations is continuous, then it is a weakly dominant strategy for an uninformed and a patient seller to set a secret reserve price s equal to v. There is an infinite number of mixed-strategy equilibria in which an uninformed and patient seller sets a secret reserve price s to v with a non-zero probability. There is no pure strategy equilibrium where the seller sets the secret reserve price s above v with probability 1, because in this case no bidder enters in equilibrium.

Since the seller does not know v, she should set a secret reserve price equal the sample analogue of v, which is the maximum of the highest order statistic. ■

We need to emphasize several important premises and consequences of Proposition 2.2. First of all, an uninformed seller should set a secret reserve price equal to the maximum of the highest order statistic only if the cost of relisting in terms of time and resources is zero. If the seller faces a sizeable cost of relisting, then there is the usual trade-off between acquisition of an additional information and foregoing profits from selling earlier. If this is the case, a seller is better off lowering a secret reserve price to capture expected profits from selling an item earlier.

Second, under assumptions of Proposition 2.2, an uninformed seller should set a secret reserve price to the maximum unsuccessful bid across all unsuccessful prior auctions. This result holds independently of whether bidders engage in sniping or not. Even if some bidders are not able to place their highest bids with a positive probability due to sniping, in the absence of any other information about the distribution function of bidders' valuations the highest unsuccessful bid across all unsuccessful auctions is the most accurate statistic of the upper bound of the support of valuations. Further, note that the maximum unsuccessful bid across all unsuccessful prior auctions is a random variable with a support $[r, v]$, where r is a public reserve price and v is the upper bound of the support of the distribution function of bidders' valuations. Hence, by part (a) Proposition 2.1 in equilibrium a seller should observe an non-zero entry in an eBay auction with a secret reserve price.

Third, the seller's choice of an optimal BIN depends on her attitude toward risk and whether a BIN is used together with a secret reserve price or not. If the BIN is used without a secret reserve price, then by rules of eBay the BIN is essentially a temporary posted price, and according to Mathews (2003), a risk-neutral seller should set the BIN equal to v, while a risk-averse seller should set the BIN equal to some value below v[21]. Under the assumption that the seller has no knowledege of the upper bound of the support v, we assume that the seller should use the maximum of the highest order statistic as a proxy for v.

If the BIN is used together with a secret reserve price, then the BIN essentially turns into a permanent posted price, since by rules of eBay, whenever a BIN and a secret reserve price are used together, the BIN is active as long as the secret reserve price is not met. According to Hidvegi et. al. (2006), in this case if either the seller or buyers are risk-averse, the seller should set the BIN at least to the bidders' expected highest valuation.

Last, when a BIN and a secret reserve price are used together, there is an equilibrium in which a BIN and a secret reserve price are simultaneously set to the upper bound of the support of the distribution function of bidders' valuations or to its sample analogue. In this case bidders can infer the size of a secret reserve prices by observing the BIN. As a result, a secret reserve price is no longer a random variable from the standpoint of bidders and entry strategies of bidders should change. However, by rules of an eBay auction, a seller can change a secret reserve price at any time prior to the 12 hours of an auction closing time. This gives bidders an incentive to participate in an auction with a secret reserve price and a BIN even if their valuations are below the BIN. We test whether bidders infer sizes of secret reserve prices from BINs by estimating the effect of a simultaneous presence of a BIN and a secret reserve price on entry. We find that the simultaneous use of BINs and secret reserve prices has a small though statistically significant negative effect on entry[22].

[21]Note that if the BIN is used without a secret reserve price, it disappears once the first bid is made.

[22]For the results of this test see Table A2 in the appendix.

2.3 Estimation Strategy and Identification Problems

To analyze seller behavior in eBay auctions, we construct determinants of public and secret reserve prices from distributions of bids, which sellers observe in eBay auctions. We assume that sellers form their beliefs about distribution functions of bidders' valuations by observing bids in eBay auctions of similar items or from unsuccessful eBay auctions for their own items[23]. To avoid making assumptions about sellers' beliefs of distribution functions of bidders' valuations, we use a non-parametric density estimator with a normal kernel to estimate distribution functions of bidders' valuations.

To resolve the issue with the lack of data given the non-parametric approach we make the following assumptions. First, we reduce the number of dimensions according to which bidders form their valuations of a tractor down to just three. These three dimensions are an age of a tractor, a horse power of a tractor, and a brand of a tractor. Second, we discretize each dimension to allow for a sufficient number of observations for a non-parametric estimation. In particular, we allow only three categories of age: 0-10 years old, 11-30 years old, and above 30 years old. Similarly, we allow only three categories of horse power: up to 20 HP, 21-40 HP, above 40 HP, and we allow only two categories of brand: whether a tractor is manufactured by John Deere or not. Thus, totally we have 18 categories of tractors.

The discretization of tractors in this way is driven by the types of tractors in each category. The division by age reflects whether a tractor on sale is of the current generation, the past two or three generations, or is an antique tractor.

The division of tractors by horse power captures the intended purpose of tractors. Tractors up to 20 HP are smaller lawn tractors such as mowers, while tractors with 21-40 HP are compact utility tractors used for everyday farming needs. Tractors with horse power above 40 HP are utility tractors used for heavier tasks such as cropping or construction.

[23] Einav et. al. (2012) make a similar argument - they argue that many sellers use online markets not only to sell their items per se, but also for experimenting with an optimal selling strategy by relisting their items multiple times with varying auction parameters or sale formats.

The division of tractors by John Deere brand and non-John Deere brands is driven by the popularity of John Deere tractors among farmers and the large number of John Deere tractors in the data.

To obtain an untruncated distribution of valuations we use bids only from auctions with a public reserve price of at most $100[24]. In addition, we ignore multi-stage auctions where there are two or more identical bidders participating in more than one stage of a multi-stage auction. We do so to remove the strategic component from the analysis, since repeated bidders and a seller may engage in a strategic dynamic behavior[25]. To account for the censoring of the highest valuations in successful auctions we run a two-step procedure[26]. First, we non-parametrically estimate distribution functions of bidders' valuations without accounting for the censoring problem. Next, we replace highest bids in successful auctions with expected valuations calculated by using non-parametrically estimated distribution functions from the first step[27]. Given the updated highest bids, we non-parametrically re-estimate distribution functions.

We use the resulting estimated distribution functions for each tractor category to calculate determinants of auction parameters set by sellers in eBay auctions. It is important to emphasize that these estimated distribution functions are not unbiased estimates of true distributions of valuations which bidders have, and the question of whether sellers are able to control for the bias at least partially is left open for a future study. Nevertheless,

[24]We do not consider auctions without high public reserve prices, because sellers most likely ignore such auctions as an accurate source of information about distribution functions of bidders' valuations given that such auctions truncate bids and act rather as posted price sales.

[25]We recognize that by ignoring auctions with repeated bidders we may overlook an important component from the analysis. However, we leave the analysis of the repeated bidder and seller interaction for a follow up paper.

[26]Note that by proposition 1, in a successful auction with sniping the winning bidder never reveals her true valuation and bids the second highest valuation or some value below it, while all other bidders reveal their valuations or some values below them.

[27]We use the standard expression for an expectation of a truncated variable to calculate an expected highest valuation in a successful auction: $E[v] = Prob(v \leq b) * b + Prob(v > b) * E[v|v > b]$, where $E[v]$ is the expected highest valuation in a successful auction, b is the highest bid in a successful auction, $Prob(v \leq b)$ is the non-parametrically estimated probability of bidders' valuations, and $E[v|v > b]$ is the expectation of the highest valuation conditional on exceeding the highest bid in a successful auction.

in the appendix we test for major determinants of the bias. We find that the presence of secret reserve prices reduces entry of bidders and particularly of low-valuation bidders. As a consequence, there is an upward bias due to under-representation of valuations in the left tail of distributions of valuations, which sellers should observe in auctions with secret reserve prices. Second, we find that bidders' valuations are not independently distributed and contain a commonly distributed component. As a result, the highest bids in eBay auctions overstate their underlying valuations, which also contributes to the upward bias. Lastly, due to sniping, realized unsuccessful bids do not always reflect underlying valuations, since some bidders may not be able to place their bids before auction closing times. Hence, the presence of sniping results in the downward bias. We leave the question of which form of the bias dominates the distributions of valuations obtained by sellers from eBay auctions for a future work.

2.4 Determinants of Auction Parameters

2.4.1 Determinants of Public Reserve Prices

Myerson (1981) and Riley and Samuelson (1981) show that when bidders' valuations are independently distributed, it is always optimal for a seller with a positive valuation of an item on sale to set a non-zero public reserve price and exclude some low-valuation bidders. The authors further show that the optimal public reserve price should satisfy the following closed form expression: $r_i^* = x_i + \frac{1}{\lambda(r_i^*)}$, where r_i^* is the optimal public reserve price for an item i, x_i is the seller's valuation of the item on sale, and $\lambda(r_i^*)$ is the hazard rate associated with the distribution function of bidders' valuations and evaluated at the seller's optimal public reserve price. When bidders' valuations have a commonly distributed component, Cai, Riley and Ye (2007) show that public reserve prices can be used as a credible signal of the value of the item on sale.

The test of an inter-dependent versus independent distribution of valuations in the appendix suggests that bidders' valuations have a commonly distributed component. Hence, public reserve prices in auctions for used tractors are used both to exclude low-valuation bidders and to credibly signal values of tractors on sale.

If public reserve prices are used together with secret reserve prices and if uninformed sellers use secret reserve prices to learn distribution functions of bidders' valuations, it is hard to justify the use of public reserve prices as a screening device. However, if public reserve prices can serve as a credible signal of quality, then uninformed sellers can obtain a more accurate right tail of distribution of valuations by setting high public reserve prices and attracting higher-valuation bidders. Next, we explicitly state the hypotheses about functions of public reserve prices in eBay auctions for used tractors.

Hypothesis 1 (public): *If bidders' valuations have both an independently and a commonly distributed component and if sellers use eBay auctions to sell, then the sellers employ public reserve prices to screen out low valuation bidders and to credibly signal the value of their items on sale.*

Hypothesis 2 (public): *If bidders' valuations have both an independently and a commonly distributed component and if sellers use auctions to gather information rather than sell, then the sellers use public reserve prices exclusively to credibly signal the value of their items on sale.*

In our test of functions of a public reserve price we include determinants accounting for the screening effect and the signalling effect. To account for the screening effect we include one over the hazard rate estimated at the public reserve price, where the hazard rate is calculated from non-parametrically estimated distribution functions of bidders' valuations. To account for the signalling effect we include the average bid among the determinants of a public reserve price. Since there is a reasonable suspicion that the public reserve price may lower the average bid by limiting entry of low valuation bidders, and our estimates may suffer from a simultaneity bias, in our estimates we use the maximum bid as an instrument for the average bid. In addition, we include tractor characteristics such as tractor age, horse power and brand to capture sellers' valuations of tractors on sale.

To test whether public reserve prices are used exclusively for signalling when sellers use public and secret reserve price at the same time, we interact signalling and screening determinants of public reserve prices with a dummy for a secret reserve price. The results of our test are presented in Table 2.7.

Variable	IV		Fixed Effects IV	
$\frac{1}{\text{Hazard rate}_j(\text{r}_i)}$*(1-Secret res. dummy$_i$)	0.172**	(0.088)	0.283***	(0.031)
Average bid$_i$*(1-Secret res. dummy$_i$)	0.257**	(0.107)	0.131***	(0.022)
$\frac{1}{\text{Hazard rate}_j(\text{r}_i)}$*Secret res. dummy$_i$	-0.432***	(0.039)	-0.333***	(0.022)
Average bid$_i$*Secret res. dummy$_i$	0.563***	(0.028)	0.501***	(0.014)
Engine HP$_i$	15.834***	(2.446)	17.821***	(1.903)
Age$_i$	-60.628***	(7.915)	-58.539***	(4.436)
John Deere$_i$	918.560***	(134.717)	890.793***	(122.691)
Constant	2516.408***	(337.854)	2480.307***	(126.929)
Overall R^2	0.644		0.609	
Number of groups			4499	
Number of observations	8268		8268	

Notes: The dependent variable is the public reserve price above $100;
Max. bid$_i$*(1-Secr. res. dummy$_i$) and Max. bid$_i$*Secr. res. dummy$_i$ are used as instruments for Aver. bid$_i$*(1-Secr. res. dummy$_i$) and Aver. bid$_i$*Secr. res. dummy$_i$;
Observations are clustered by seller ID in FE model; Robust standard errors in parentheses; ***, **-statistical significance at 1%, 5%;

Table 2.7. Determinants of public reserve prices

In both models in Table 2.7, when secret reserve prices are not used, the determinants of a public reserve price accounting for screening and signalling effects are positive and statistically significant. However, when public reserve prices are used together with secret reserve prices, the coefficient for one over the hazard rate becomes negative, violating the screening assumption, while the determinant for a signalling effect in the fixed effects IV model more than doubles in size in comparison to the case without a secret reserve price. We interpret this result as an evidence that when sellers use auctions to gather information about bidders' valuations and employ secret reserve prices, the sellers use public reserve prices exclusively to signal the quality of their tractors on sale and to obtain valuations of their tractors primarily from higher-valuation bidders. However, when sellers use auctions exclusively to sell their items and do not empoy secret reserve prices, the sellers use public reserve prices both to screen out low valuation bidders and to signal the quality of their items.

2.4.2 Determinants of Secret Reserve Prices

To test the purpose of a secret reserve price, we suggest three possible hypotheses. Under the first hypothesis, a secret reserve price is an alternative to a public reserve price and is set to screen out bidders with valuations below the threshold level indicated by the size of a secret reserve price[28]. Under this hypothesis a secret reserve price is a function of one over the hazard rate estimated at the secret reserve price and a seller's valuation of the item on sale[29]. The main argument against this hypothesis is the prevalence of a simultaneous use of public and secret reserve prices in eBay auctions.

Hypothesis 1 (secret): *If bidders' valuations have both an independently and a commonly distributed component, sellers use secret reserve prices exclusively to screen out low valuation bidders.*

[28]The view of a secret reserve price as an alternative to a public reserve prices has received the most attention in the literature (see, for example, Katkar and Reiley (2006), Lucking-Reiley et. al. (2007) for empirical tests).

[29]Under this hypothesis a secret reserve price cannot serve the signalling function of a public reserve price, because bidders do not observe it.

Under the second hypothesis a very patient informed seller uses a secret reserve price to sell at the highest possible price. In this case an informed seller sets a secret reserve price at the known upper bound of the support of the distribution function hoping that the highest-valuation bidder enters an auction. Under this hypothesis, an informed patient seller uses an auction as a posted price sale where the sale price is revealed through bidding.

There are two main arguments against this hypothesis. The first one is the prevalence and stability of the share of unsold items in eBay auctions. If sellers had a perfect knowledge of the distribution function of bidders' valuations and were sufficiently patient, then we would observe an increase in the share of sold tractors with the number of relistings. However, we see that the share of sold tractors of about 30% is relatively stable across relistings. The second argument against this hypothesis is the availability of a cheaper posted price sale format on eBay, which an informed seller should prefer to a more costly auction format with a secret reserve price[30].

Hypothesis 2 (secret): *If bidders' valuations have both an independently and a commonly distributed component, informed and patient sellers use secret reserve prices to sell at the highest possible price.*

Under the third hypothesis, a patient uninformed seller uses a secret reserve price to learn parameters of the distribution functions of bidders' valuations and to sell an item at the highest possible price given the knowledge of these parameters. In this case the equilibrium secret reserve price is set at the seller's belief about the upper bound of the support of the distribution function of bidders' valuations. As the uninformed patient seller observes unsuccessful bids, she updates her belief about the support and other parameters of the distribution function and recalculates the secret reserve price. By Proposition 2.2 the seller's strategy under this hypothesis is to set the secret reserve price at the max-

[30]The cost of an unsuccessful posted price sale on eBay is \$0.50 as opposed to a US 25.25 dollar cost of running an unsuccessful auction with a secret reserve price and a BIN. Of course, this argument holds under the assumption that the pool of buyers in auctions and posted price sales is the same.

imum of the observed highest order statistic or at the highest bid across all unsuccessful previous auctions.

Hypothesis 3 (secret): *If bidders' valuations have both an independently and a commonly distributed component, uninformed and patient sellers use secret reserve prices to learn parameters of the distribution function bidders' valuations.*

First, we test hypothesis 1 against hypotheses 2 and 3. Under all three hypotheses, an introduction of a secret reserve price should lower the likelihood of a successful sale. However, under hypothesis 1, the negative effect of a secret reserve price on the probability of sale should be of a comparable size as the effect of a public reserve price. The regression results in Table 2.8 show that a dollar increase in the public reserve price has a small negative and statistically insignificant effect on the probability of sale, while a dollar increase in the secret reserve price has a much stronger negative and statistically significant effect on the probability of sale.

In the next test we explicitly test possible determinants of secret reserve prices. In our sample we have data on secret reserve prices only in unsuccessful auctions. In successful auctions the data on secret reserve prices is replaced by sale prices. Further, the secret reserve prices in our data are the ones that were active in the last 12 hours of auction closing times. Hence, the secret reserve prices in our data do not necessarily equal the secret reserve prices set initially at auction start times[31]. To account for the missing secret reserve prices in successful auctions, we use average secret reserve prices calculated for each of the 18 categories of tractors.

[31]Note that by eBay auction rules a seller can update a secret reserve price at any time before the last 12 hours of an auction closing time.

Variable	Probit Marginal Effect	Standard Error
Secret reserve$_i$	-0.00005***	0.000002
Public reserve$_i$	-0.000001	0.000005
BIN price$_i$	-0.000001	0.000001
John Deere$_i$	0.019	0.015
Age$_i$	0.0007	0.0005
Engine HP$_i$	-0.0004*	0.0002
Seller feedback score$_i$	0.00005*	0.00002
Number of bidders$_i$	0.021***	0.002
Average bid$_i$	0.00003***	0.00001
Number of observations		13051

Notes: The dependent variable is Sale (Y/N); Standard errors are clustered by Seller ID; ***,*-statistical significance at 1%, 10%;

Table 2.8. Determinants of a successful sale (pooled data from all auctions)

To construct a sample for the test, we pool together tractors from all relistings. We ignore single-stage auctions with secret reserve prices, because some single-stage auctions are relistings of tractors listed before the data was collected and whose determinants we do not observe.

Lastly, since our test of determinants of secret reserve prices is based on the assumption that sellers do not have knowledge about distribution functions of bidders' valuations, it is necessary to specify initial beliefs of sellers about these distribution functions. We assume the following process of formation of beliefs. Initially, uninformed patient sellers set their secret reserve prices by observing highest bids in auctions of similar items. Once they are able to run their own unsuccessful auctions and gather information about valuations of their own tractors, they update their secret reserve prices by incorporating the data from unsuccessful auctions. To account for the initial beliefs of sellers we average highest bids from auctions in each of the 18 categories of tractors and include these values among determinants of secret reserve prices.

In our test of determinants of secret reserve prices we include the following regressors: $\{\widehat{b}_{i,t-k}, \overline{\overline{b}}_j, \frac{1}{\lambda_j(s_{i,t})}, x_i, m_i\}$, where $\widehat{b}_{i,t-k}$ is the highest bid from all previous unsuccessful auctions for tractor i with the highest bid made during an unsuccessful auction listing $k \in \{1,2,3,4\}$, $\overline{\overline{b}}_j$ is the average highest bid for a tractor in category j, $\frac{1}{\lambda_j(s_{i,t})}$ is one over the estimated hazard rate for a tractor in category j and calculated at the secret reserve price $s_{i,t}$, x_i is the set of parameters accounting for tractor heterogeneity such as tractor's age, horse power and brand, and $m_i \in \{1,2,3,4\}$ is the number of times tractor i was previously listed.

The results of both models in Table 2.9 allow us to reject hypothesis 1 that sellers use secret reserve prices to screen out low valuation bidders. In the OLS model, the screening effect represented by one over the hazard rate is negative and statistically insignificant, while in the fixed effects model the screening effect is statistically significant but negative, which violates the screening assumption.

Further, in both models by far the most significant determinants of secret reserve prices are the highest bids from previous unsuccessful auctions, which capture sellers' acquisition of information from running unsuccessful eBay auctions for their individual items, and average highest bids for tractors in a similar category, which capture sellers' acquisition of information from observing eBay auctions for similar items.

To see that secret reserve prices are used by uninformed or partly informed sellers for the purposes of learning, note that in Table 2.9 the coefficient for the number of relistings is statistically significant and negative. In other words, there is a downward adjustment in secret reserve prices with an additional relisting. This contradicts hypothesis 2 that informed sellers use secret reserve prices solely to sell at the highest possible prices, since such informed sellers should not adjust secret reserve prices based on information acquired from running additional unsuccessful relistings.

Variable	OLS		Fixed Effects	
Highest $\text{Bid}_{i,t-k}$	0.589***	(0.058)	0.521***	(0.133)
Average Highest Bid_j	0.553***	(0.119)	0.576***	(0.197)
$\frac{1}{\text{Hazard rate}_j(s_{i,t})}$	-0.145	(0.136)	-0.463**	(0.204)
Age_i	-16.007**	(7.358)	-78.649**	(32.886)
Engine HP_i	17.961***	(6.754)	42.246*	(22.774)
John Deere_i	1300.824***	(342.253)	2297.449**	(904.939)
Number of relistings_i	-380.236**	(154.887)	-349.718**	(158.671)
Constant	-986.682**	(477.178)	374.094	(750.065)
Overall R^2	0.759		0.738	
Number of groups			489	
Number of observations	782		782	

Notes: The dependent variable is secret reserve price $s_{i,t}$; Robust standard errors in parentheses; Observations and errors are clustered by seller ID in FE model; ***,**,* - statistical significance at 1%, 5%, 10%;

Table 2.9. Determinants of secret reserve prices

To explicitly test hypothesis 2 that informed sellers use secret reserve prices exclusively to sell at the highest possible prices against hypothesis 3 that uninformed sellers use secret reserve prices to learn, we regress current secret reserve prices on current and past highest bids. If sellers are informed, then the variation in the current secret reserve price should be better explained by the current highest bid then by the past highest bid. However, if sellers are uninformed, then the variation in the current secret reserve price should be better explained by the past highest bid.

Variable	OLS Coefficient		OLS Coefficient	
Highest Bid$_{i,t}$	0.013	(0.008)	0.017	(0.019)
Highest Bid$_{i,t-k}$	0.587***	(0.057)		
Average Highest Bid$_j$	0.445***	(0.092)	1.001***	(0.080)
Number of relistings$_i$	-369.971**	(153.222)	37.384	(222.001)
Age$_i$	-12.466*	(7.015)	-68.589***	(8.793)
Engine HP$_i$	15.409**	(6.733)	51.711***	(6.862)
John Deere$_i$	1164.441***	(352.464)	2534.878***	(377.694)
Constant	-814.972*	(477.040)	-1419.283**	(598.874)
R^2	0.761		0.591	
Number of observations	782		782	

Notes: The dependent variable is secret reserve price$_{i,t}$; Robust standard errors in parentheses; ***,**,*-statistical significance at 1%, 5%, 10%;

Table 2.10. Informed versus uninformed sellers

The variation in the current secret reserve price in Table 2.10 is better explained by the past highest bid than by the current highest bid. In fact, in both models in Table 2.10, the current highest bid is statistically insignificant. Hence, there is more evidence to accept hypothesis 3 than hypothesis 2. However, since there is a strong correlation between highest unsuccessful bids across different periods, it is hard to make a definite claim about exact determinants of a secret reserve price. In addition, since initially uninformed sellers become more informed once they run unsuccessful auctions or observe bidding behavior in auctions for similar items, it is virtually impossible to separate the behavior of a completely informed seller from the behavior of a partially informed seller by observing only the variation in secret reserve prices. Guided by these considerations, we do not make a definite claim that secret reserve prices are used exclusively to learn and that sellers who use secret reserve prices are completely uninformed.

2.4.3 Determinants of BINs

To test determinants of BINs, we invoke theoretical results of Mathews (2003) and Hidvegi et. al. (2006). The size of a BIN in an eBay auction depends on whether the BIN is permanent or temporary. According to Mathews (2003), a risk-neutral seller should set a temporary BIN equal to the upper bound of the support of the distribution function of bidders' valuations, while a risk-averse seller should set a temporary BIN to some value below the upper bound. According to Hidvegi et. al. (2006), a risk-neutral and a risk-averse seller should set a permanent BIN at least to the expected highest valuation. Since the upper bound of the support of valuations is strictly above the expected highest valuation, we conclude that a permanent or a temporary BIN in an eBay auction should be bounded above by the upper bound of the support of valuations and bounded below by the expected highest valuation.

Under the assumption that a seller does not have perfect information about the distribution functions of bidders' valuations, and therefore, cannot set the temporary BIN to the upper bound of the support of bidders' valuations, we assume that such an uninformed seller uses the highest bid from previous unsuccessful auctions to determine the upper bound. To account for the expected highest valuation, we use an average highest bid for each of the 18 tractor categories. Further, we interact these determinants with secret reserve prices to see whether sellers set BINs differently in auctions with secret reserve prices.

To test the determinants of BINs, we pool together auctions with 2 or more listings. We ignore single-listing auctions, because for these auctions we do not observe the highest unsuccessful bids from previous listings. We presents the results of our test in Table 2.11.

Variable	OLS	Fixed Effects
Highest Bid$_{i,t-k}$	1.013*** (0.032)	0.715** (0.358)
Average Highest Bid$_j$	0.151*** (0.058)	0.256 (0.389)
Highest Bid$_{i,t-k}$*Secret reserve$_i$	0.123*** (0.046)	0.368 (0.345)
Average Highest Bid$_j$*Secret reserve$_i$	-0.042 (0.059)	-0.232 (0.346)
Constant	185.674 (266.492)	1183.553 (822.096)
Overall R^2	0.936	0.933
Number of groups		219
Number of observations	354	354

Notes: The dependent variable is BIN$_{i,t}$; Robust standard errors in parentheses; Observations and errors are clustered by seller ID in FE model; ***,** - statistical significance at 1%, 5%;

Table 2.11. Determinants of BINs

The results in Table 2.11 show that independently of whether a secret reserve price is used or not, by far the most important determinant of a BIN in terms of size and significance is the highest bid from previous unsuccessful auctions. Hence, we conclude that in eBay auctions sellers tend to set BINs closer to their beliefs about upper bounds of bidders' valuations rather than to their beliefs about expected highest valuations.

It is important to emphasize that in our tests of determinants of auction parameters the determinants of BINs and secret reserve prices are identical. We show that the variation in both BINs and secret reserve prices is largely explained by the variation in previous highest bids and expected highest valuations. The main difference is in the weights of these two determinants. While the variation in secret reserve prices is explained almost equally by the variation in previous highest bids and expected highest valuations, the variation in BINs is almost exclusively determined by the variation in the previous highest bids only. Given these findings, we conclude that the data on BINs can serve as a relatively accurate approximation of the data on secret reserve price, whenever the data on secret reserve prices is not available or impossible to obtain.

2.5 Determinants of Sale Prices

To test the effect of secret reserve prices on sale prices, we pool together tractors which were sold after up to 5 relistings on eBay. We test sale effects of reserve prices only for those tractors which were sold on eBay and discard those tractors which were listed but never sold on eBay. We do so, because we do not observe final selling prices of unsold tractors as sellers of such tractors might be using other selling formats or other selling platforms after delisting from eBay. We also omit observations with more than 5 relistings, because sellers of such tractors probably use the eBay platform for purposes other than selling or learning bidders' valuations [32].

To test sale effects of a secret reserve price, we decompose the total sale effect of a secret reserve price into an information-acquisition effect and any other possible effect. To capture the information-acquisition effect of a secret reserve price, we include among the

[32]We find that several online stores use eBay as an advertising outlet by constantly relisting an item and providing links to their websites in item description areas.

regressors the total number of times a secret reserve price has been used. If secret reserve prices are used to acquire information about bidders' valuations, then the frequency of the use of a secret reserve price should serve as an indication of the degree of seller knowledge about the distribution function of bidders' valuations. To capture any other possible sale price effect of a secret reserve price, we include a secret reserve price at the time of a successful auction listing. In addition, among the determinants of sale prices we include the total number of relistings, a public reserve price at sale, a BIN at sale, and tractors characteristics such as age, horse power and brand categories.

Table 2.12 shows that the non-information-acquisition effect of a secret reserve prices, indicated by the coefficient of a secret reserve price at sale, is positive, statistically significant, and comparable in scale to positive sale effects of a public reserve price and a BIN. The information-acquisition effect of a secret reserve price, indicated by the number of times a secret reserve price has been used, also has a sizeable positive and a statistically significant sale effect. Hence, contrary to previous empirical studies of secret reserve prices, we provide statistical evidence that the use of a secret reserve price has a strong positive effect on a sale price[33].

We believe that previous studies have different results, because they test the impact of secret reserve prices under the assumption that sellers do not relist items after an unsuccessful sale. Since the presence of a secret reserve price is highly correlated with the probability of relisting, in the absence of an explicit control for relistings the positive impact of a secret reserve price on a sale price is neutralized by the negative effect of relistings. In fact, if we sum up coefficients of the number of times a secret reserve price has been used and the number of relistings in Table 2.12, the resulting small positive difference becomes statistically insignificant[34].

To make sure that there is no endogeneity problem and the use of secret reserve prices

[33]See, for example, Bajari and Hortacsu (2003) and Katkar and Reiley (2006).

[34]The difference between these two coefficients in OLS model has an F-statistic of 0.55 and a p-value of 0.459. In FE model, the difference has an F-statistic of 0.39 and a p-value of 0.534.

is uncorrelated with an unobserved heterogeneity in the quality of tractors, we regress OLS residuals from Table 2.12 on the set of regressors in Table 2.12. We find that residuals are not correlated with any of the regressors[35].

Variable	OLS	Fixed Effects
[Number of times a secret reserve price is used]$_i$	0.120*** (0.022)	0.103*** (0.034)
[Number of relistings]$_i$	-0.106*** (0.029)	-0.087** (0.041)
[Secret reserve at sale]$_i$	0.00002*** (0.000003)	0.00001*** (0.000005)
[Public reserve at sale]$_i$	0.00005*** (0.000002)	0.00005*** (0.000007)
[BIN at sale]$_i$	0.00002*** (0.000001)	0.00002*** (0.000004)
Age$_i$	-0.027*** (0.0007)	-0.025*** (0.002)
Engine HP$_i$	0.008*** (0.0003)	0.009*** (0.001)
Constant	8.744*** (0.040)	8.629*** (0.076)
Overall R^2	0.509	0.502
Number of groups		3942
Number of observations	8316	8316

Notes: The dependent variable is log(sale price$_i$); Robust standard errors in parantheses; Observations and errors are clustered by seller ID in FE model; ***-statistical significance at 1%; We do not report tractor brand dummies; Tractor brand dummies are jointly statistically significant at 1% with p-value<0.0001;

Table 2.12. Determinants of sale prices

[35] We run an additional endogeneity test by regressing a secret reserve dummy on tractor characteristics such as tractor age, horse power, engine hours and eleven brand categories. We find that tractor age and horse power are the only statistically significant determinants of the use of a secret reserve price. The rest of tractor characteristics are not statistically significant suggesting that there is no unobserved tractor heterogeneity correlated with the use of secret reserve prices.

To get a better idea about the impact of secret reserve prices on sale prices, we run a simulation exercise where we implement different strategies of an uninformed seller. To introduce dynamics into seller actions, we impose a cost of an additional relisting. We introduce a relisting cost, because by Proposition 2.2, an uninformed seller who does not face any costs of relisting and cares only about learning the distribution function of bidders' valuations does not have an incentive to adjust a secret reserve price. In the presence of relisting costs, the seller has to balance the expected payoff from keeping a secret reserve high and learning bidders' valuations against the expected payoff from selling an item earlier and avoiding the relisting costs.

In our simulation exercise we impose three kinds of relisting costs. The first cost is the listing fee of $25. The second one is a future discount factor of 0.99. The third cost is that the seller's payoff becomes zero if she does not sell an item after the fifth relisting. To make our simulation exercise specific, we assume that a seller wants to sell a John Deere tractor with a 95 HP engine and which is about 37 years old. There are 25 of these tractors with a public reserve price less than US 100 dollars. We use 104 bids from these 25 tractors to non-parametrically estimate a cumulative distribution function of bidders' valuations and its associated inverse cumulative distribution function. We draw five valuations from the estimated inverse cdf and calculate the seller's revenue in an ascending second-price auction under different regimes.

We assume that the seller is uninformed about the estimated distribution function and only knows the maximum sale price of John Deere tractors with ages above 30 and horse power above 40. The maximum sale price of these tractors in our sample is $18988. Under the first regime, the seller sells an item in an ascending second-price auction without a public or a secret reserve price. As a result, she does not need to invoke her beliefs about the parameters of the distribution function and always sells at the second highest valuation at the first listing.

Under the second regime, the seller uses a secret reserve price. Since the seller is uninformed, it is optimal to set a secret reserve price equal to the expected maximum of the highest order statistic. In our case it is \$18988 - the maximum sale price of John Deere tractors with ages above 30 and HP above 40. Once the seller is able to observe the highest unsuccessful bid from the an auction for her specific John Deere tractor, the seller can adjust the secret reserve price accordingly. We assume that the seller follows a simple updating rule of the following form, $S_t = \alpha * S_{t-1} + (1-\alpha) * Y^1_{t-1}$, where S_t is a secret reserve at time t, S_{t-1} is the previous period secret reserve, Y^1_{t-1} is the highest unsuccessful bid in the previous auction stage, and α, $\alpha \in [0,1]$, is the weight the seller attaches to the previous period secret reserve price.

After the seller sets the initial secret reserve price, we calculate the seller revenue in an ascending second-price auction given the listing cost of \$25 and the draw of 5 independent valuations from the estimated distribution of valuations for the 37 year old John Deere tractors with 95 HP. If the highest valuation is below the initial secret reserve price, we proceed to the next stage, where we draw another set of 5 valuations, subtract \$25 listing fee, and discount the seller revenue at the rate of 0.99. We continue drawing new sets of 5 valuations, charging \$25 listing fee, and discounting seller payoff if there is no sale up to 5 times. If there is no sale at the fifth draw, the seller obtains zero payoff. We run 10000 iterations to obtain an average seller revenue for 100 weights α, where α ranges from 0 to 1 with an increment of 0.01. The simulation results are presented in Table 2.13.

The results in the table show that depending on the updating rule, the introduction of a secret reserve price can raise the simulated mean price quite significantly, however, at the cost of an additional relisting. For example, the simulated mean seller revenue with $\alpha = 0$, where the seller adopts the highest unsuccessful bid as a secret reserve price and completely ignores the prior secret reserve price, is \$1547.7 more than the benchmark case without a public or a secret reserve price.

Auction Format and Data Type	Mean Price in dollars	Standard Deviation	Mean Number of Relistings
Actual data	7947.1	3357.7	1.08
Simulated without Public or Secret Reserve	8333.3	2421.6	1.00
Simulated with Secret Reserve and $\alpha = 0$	9881.4	1795.5	2.70
Simulated with Secret Reserve and $\alpha = 0.15$	9979.5	2141.8	2.90
Simulated with Secret Reserve and $\alpha = 0.5$	8311.2	4773.5	4.50
Simulated with Secret Reserve and $\alpha = 0.75$	1361.3	3998.5	>5
Simulated with Secret Reserve and $\alpha = 1$	-125.0	0.0	>5

Table 2.13. Simulated mean price of a 37 year old John Deere tractor with 95 HP

In addition, note that when $\alpha = 0$ and the seller sets the secret reserve price equal to the highest unsuccessful bid from the first stage, the seller does not need to relist the unsuccessfully sold item in an auction in the second stage. Instead, the seller can switch to a posted price sale with a price equal to the highest bid from the unsuccessful first-stage auction. Under this strategy that the seller sets a secret reserve price at the maximum of the highest order statistic, the seller's revenue from a posted price sale is exactly equal to the seller revenue from a repeated auction with a secret reserve price and $\alpha = 0$. Hence, if sellers in general have weight $\alpha = 0$ and the cost of a posted price sale is lower than that of an auction, it is optimal for sellers to switch to a posted price sale after running just one unsuccessful ascending second-price auction[36]. This can explain a large number of delisted tractors after one or two unsuccessful auction rounds in the data.

[36]Recall that eBay offers a posted price sale format in addition to an auction-type sale format.

The results in our analysis crucially depend on the assumption that entry is exogenous and that bidders' entry decisions do not depend on the presence of a secret reserve price. However, since the higher revenue in an auction with a secret reserve price is driven by the ability of a seller to sell to the highest valuation bidder at the highest possible price and entry decisions of highest valuation bidders are least affected by the presence of a secret reserve price, we expect that the lack of exogeneity of the number of bidders in auctions with secret reserve prices should not bias our results too much.

Another critical assumption is that we are assuming that bidders do not act strategically in repeated auctions and do not shade their bids. This assumption is satisfied for large sale platforms such as eBay, where many sales are conducted at the same time and where it is difficult to coordinate participation in repeated auctions[37]. In live auctions, or in electronic auctions with costly or restricted entry, the assumption of non-strategic bidders is harder to justify. However, if it is the case that bidders are strategic, then we should expect the convergence of sale prices in repeated auctions to sale prices in static auctions as was shown by McAffee and Vincent (1997).

Nevertheless, the results of the simulation exercise show that the use of a secret reserve price can potentially raise sale prices if sellers are sufficiently patient and do not face high costs of relisting, while bidders bid according to strategies of a single-stage ascending second-price auction and do not take into account the possibility of a future resale. The cost of this increase in seller revenue is the necessity to run at least one unsuccessful ascending second-price auction.

[37]We do not consider sellers' response to bidders' strategic behavior in repeated auctions, because in eBay auction sellers cannot differentiate bidders' identities and, therefore, most likely treat all bidders as non-strategic.

2.6 Concluding Remarks

In this chapter we attempt to explain several stylized facts found in the eBay data such as the widespread use of secret reserve price, the presence of relisted items and a large number of unsuccessful sales. We believe that sellers relist their items because they want to learn bidders' valuations. We also believe that there are many unsuccessful auctions on eBay, because uninformed sellers switch to other selling formats after acquiring enough information by running unsuccessful auctions[38]. We cannot demonstrate this argument empirically, because to test this argument, we need data from alternative sales venues used by sellers who delist their items from eBay or about eBay posted price sales.

In the chapter we give an explanation to why sellers may use a secret reserve price. We argue that a secret reserve price serves as a cheap and an efficient tool for uninformed patient sellers to run unsuccessful auctions to learn parameters of the distribution function of bidders' valuations. Our findings depends on two key assumptions: that sellers' do not know parameters of the distribution function of bidders' valuations and that sellers can observe truthful revelation of bidders' valuations in unsuccessful ascending second-price auctions.

There are two directions of future empirical and theoretical research. The first direction is related to the analysis of seller behavior under non-commitment, or in the presence of strategic bidders, who anticipate that unsold items will be relisted later and who take this fact into consideration. The empirical analysis of such a model involves a structural estimation of a dynamic game with evolving state variables.

The second direction of future research involves an analysis of hybrid sales, where an uninformed seller makes a choice for an optimal selling mechanism as she acquires more and more information about the distribution function of bidders' valuations. The analysis

[38]In particular, sellers can switch to posted prices sales on eBay after running unsuccessful eBay auction sales.

of such hybrid sales requires the data on different selling formats and the degree of seller informativeness.

Chapter 3

Quality Sorting Between Electronic and Physical Platforms

In this chapter we study a model of quality sorting between electronic and physical platforms. In the model a seller can auction an item with both opaque and transparent quality attributes in either platform. Bidders can observe perfectly the quality of both the transparent and the opaque attribute in a physical platform. In an electronic platform bidders can observe perfectly only the quality of the transparent attribute but not the opaque attribute. In the model, a seller pays a commission for listing her item in a physical platform. There is no commission for listing in an electronic platform. Similarly, bidders do not pay any fees for participating in either platform. We use Spence's signalling model (1973) to derive seller's equilibrium listing strategy. We find that conditional on the quality of the transparent attribute, the quality of the opaque attribute in a physical platform is always no worse than the quality of the opaque attribute in an electronic platform. We further find that when items on sale have both transparent and opaque attributes, it is impossible to compare the item's overall quality in both platforms without restricting bidders' beliefs about the quality of the opaque attribute in an electronic platform or without introducing an additional structure into the model. The main conclusion is that it is not always true that the overall quality of items in an electronic platform is necessarily lower than the overall quality of items in a physical platform.

The available studies of unobservable quality in electronic platforms use Akerlof's (1970) model of adverse selection (e.g. Lewis (2011), Overby and Jap (2009), Wolf and Muhanna (2005)). This approach suggests that since quality is unobservable in electronic platforms, the items listed in an electronic platform are sold at a discount. This drives higher quality items to a competing physical platform, where the quality is observable and there is no price discount. As a result, the overall quality of items remaining in an electronic platform is lower than in a competing physical platform. We take a somewhat different approach, and analyze the quality sorting between electronic and physical platforms within the framework of Spence's (1973) signalling model. We view the seller's decision to list in a physical platform or in an electronic platform as an informative signal, which bidders use to update their belief about the quality of items in an electronic platform. We find that depending on the size of the updated belief, sellers of higher quality items may list both in an electronic and in a physical platform. Further, we introduce two quality dimensions and show that without restricting bidders' beliefs it is impossible to compare the overall quality of items in competing electronic and physical platforms.

Our model is related to three strands in the literature. The first strand includes models of asymmetric information. There are two basic frameworks for the analysis of models with asymmetric information: screening or adverse selection models (e.g. Mussa and Rosen (1978), Maskin and Riley (1984)) and signalling models (e.g. Spence (1973)). We use the signalling model to analyze unobservable quality in an electronic platform with seller's choice of a listing platform as a signal of quality.

The second strand of related literature is the literature on two-sided platforms (e.g. Rochet and Tirole (2003), Anderson and Coate (2005), Armstrong (2006)). The literature on two-sided platforms emphasizes the optimality of fees imposed on platform users by a platform owner. In our model we introduce only one kind of a platform fee: the commission which the seller has to pay for using the physical platform. There are no other platform fees in our model. However, our model can be easily extended if we introduce differentiated fees on bidders and sellers in electronic and physical platforms.

The third related strand involves literature on competing platforms (e.g. Ellison and Fudenberg (2003), Ellison, Fudenberg, and Mobius (2004)). The key difference between our model and the models of competing platforms is our assumption that the number of bidders per seller is the same in electronic and physical platforms, and a seller does not affect the seller-bidder ratio by choosing one platform over another. This assumption can be relaxed if we endogenize bidders' entry decisions in our model.

Lastly, our model is related to the model of quality sorting between electronic and physical platforms in Jin and Kato (2007). As opposed to the quality sorting model of Jin and Kato (2007), we first derive an equilibrium seller strategy and then derive the quality ranking between the two platforms. The results in our one-dimensional quality model are similar to the results of the one-dimensional quality sorting model of Jin and Kato (2007). However, when we introduce two-dimensional quality, our results are neither consistent with the conclusions of Jin and Kato (2007) nor with the conclusions of our one-dimensional quality model.

The main advantage of our modelling approach is the ability to extend our model in multiple directions. By introducing differentiated listing fees on sellers and buyers we can connect our model to the standard models of two-sided platforms. By endogenizing bidders' entry decisions, we can obtain predictions both on the relative sizes of competing platforms and quality comparisons across platforms.

In the first 4 sections of the chapter we analyze a simple model with a single opaque quality attribute. In section 3.5 we extend the model and introduce a second transparent quality attribute. In section 3.6 we discuss the results and in section 3.7 we conclude.

3.1 Sellers

Consider a seller $s \in S$ who wants to sell an item of opaque quality q_i, which can be high or low, $i \in \{H, L\}$. Define an item of high opaque quality as type q_H and an item of low opaque quality as type q_L. The probability that an item is of each type is determined by nature and is perfectly observable by all players in the model. In particular, we denote the probability of a type q_H item by α and the probability of a type q_L item by $1 - \alpha$ and assume that α is common knowledge. A seller has no valuation for the item itself.

A seller s can offer an item of type $q_i \in \{q_H, q_L\}$ for sale in an electronic platform or in a physical platform. The key difference between these two platforms is that buyers can perfectly observe the opaque quality parameter in a physical platform and cannot observe the opaque quality parameter in an electronic platform.

By listing an item in a physical platform, a seller pays the physical platform owner a sale fee of δp, where p is the sale price and $\delta \in (0, 1)$ is a fixed share of the sale price (commission), where δ is common knowledge. The seller's listing cost in an electronic platform is normalized to zero. The sale format in both the electronic and the physical platform is the same, and it is a second-price open outcry (English) auction[1]. In the rest of the discussion we use the terms "platform" and "auction" interchangeably.

3.2 Bidders

A seller s faces N_s potential identical bidders in each platform. Each bidder demands only one item and derives her valuation of the item of each type from a corresponding distribution function. A bidder's valuation of a type q_H item, v_H, is identically and independently distributed with a continuous cumulative distribution function $F_H(v)$ and a positive support $[0, V]$. Similarly, a bidder's valuation of a type q_L item is identically and independently distributed with a continuous cumulative distribution function $F_L(v)$ and a positive support $[0, V]$. In addition, we assume that these distribution functions are stochastically independent, and the distribution function of valuations of a type q_H item first-order stochastically dominates the distribution function of valuations of a type q_L item, or that $F_H(v) \leq F_L(v)$ for all $v \in [0, V]$.

[1]The results in the model do not rely on the auction format. By the Revenue Equivalence Theorem, when risk-neutral bidders have independently and identically distributed valuations, the expected price in a first-price auction and in a second-price auction is the same.

We assume that bidders are indifferent between participating in an electronic or in a physical auction, and bidders' valuations are determined only by their perception of an item type. We further assume that the number of potential bidders in both platforms is the same and is always equal to N_s[2].

3.3 Strategies and Payoffs

We assume that a seller of a single item has perfect knowledge of the type of her item and decides between listing her item in a physical platform or in an electronic platform[3]. Hence, the seller employs a behavioral strategy $\{(\beta, 1-\beta), (\gamma, 1-\gamma)\}$, where $\beta \in [0,1]$ is the probability of listing a type q_H item (high opaque quality) in a physical auction, $(1-\beta)$ is the probability of listing a type q_H item in an electronic auction, $\gamma \in [0,1]$ is the probability of listing a type q_L item (low opaque quality) in a physical auction, and $(1-\gamma)$ is the probability of listing a type q_L item in an electronic auction.

Since buyers have identically and independently distributed valuations, a representative bidder employs an equilibrium bidding strategy in a second-price auction, which depends on the distribution function of valuations given the item's type. According to the clock model of Milgrom and Weber (1982), in a second-price open outcry auction a bidder with identically and independently distributed valuation $v_i \sim F_i$ bids her valuation v_i unless she is the last bidder, in which case she bids the price at which the previous remaining bidder dropped out. Given this equilibrium bidding strategy, by listing an item in a platform with identical bidders who derive their valuations from the

[2]Since bidders are symmetric, the possibility that the same bidders may participate in different auctions does not affect the results in the model.

[3]We rule out the possibility that a seller may list her item in two platforms simultaneously, given that a sale in any platform is binding. Hence, it is not possible that a seller conducts two sales of a single item at the same time in different platforms and then chooses a sale with the highest realized price. While some auctions on eBay or other electronic platforms sometimes intimate that the electronic auction can be truncated by sale of the item in simultaneous physical auctions, we leave analysis of such a situation as an extension to be considered in future work.

same distribution function F_i, a seller obtains the expected price $p_i = \int_0^V v dG_i(v)$, where $G_i(v) = NF_i(v)^{N-1} - (N-1)F_i(v)^N$ is the distribution of the second-highest order statistic[4].

Note, however, that both the equilibrium bidding strategy and the expected price p_i depend on the distribution function of bidders' valuations, which in turn depends on the bidders' belief about the type of an item on sale. In particular, since bidders fully observe the opaque quality parameter of an item listed in a physical platform, for each item type in a physical auction bidders' derive their valuations from its corresponding distribution function. However, when bidders participate in an electronic platform and cannot observe the opaque quality parameter, they derive their valuations from a mixture distribution function, which is a convex combination of a distribution function corresponding to an item with high opaque quality and a distribution function corresponding to an item with low opaque quality.

Since bidders are identical, they form a common belief about the quality of an item on sale. Let $\theta \in [0,1]$ denote the common belief that an item on sale in an electronic platform is of high opaque quality and $(1-\theta)$ denote the common belief that an item on sale in an electronic platform is of low opaque quality. Then we can introduce expected prices for each platform given bidders' beliefs about item types. Since there are two distinct item types and bidders can perfectly distinguish between them in a physical platform, we define a menu of two prices in a physical platform. In an electronic platform bidders cannot distinguish between items of high and low opaque quality; hence, in an electronic platform we define only one price.

The menu of prices in a physical platform is $\{p_H^P, p_L^P\}$, where the subscript denotes quality and the superscript denotes that the price is formed in a physical platform. Each

[4]For derivation of the expected price, see Milgrom and Weber (1982) or Krishna (2009).

price in the menu is defined as follows, $p_i^P = \int_0^V v dG_i(v)$, $i \in \{H, L\}$. The expected price in an electronic platform, $p^E(\theta)$, is a function of θ, the bidders' belief that the item on sale is of high quality. The superscript "E" denotes that the price is formed in an electronic platform. Given the belief θ, we define $p^E(\theta) = \int_0^V v dG(v, \theta)$, where $G(v, \theta) = N(\theta F_H(v) + (1 - \theta) F_L(v))^{N-1} - (N-1)(\theta F_H(v) + (1-\theta) F_L(v))^N$ is the distribution of the second-highest order statistic for a mixture distribution $\theta F_H(v) + (1-\theta) F_L(v)$ for all $\theta \in [0, 1]$.

3.4 Equilibrium

Before we proceed to equilibrium predictions, in Lemma 3.1 we derive the ranking of expected prices in an electronic and a physical platform and the dependence of expected prices in an electronic platform on the belief parameter θ.

Lemma 3.1

a)For $v \in [0, V]$, if $F_H(v) \leq F_L(v)$, then $p_L^P \leq p^E(\theta) \leq p_H^P$ for $\theta \in [0, 1]$.

b)For $v \in [0, V]$, if $\theta_1 < \theta_2$, then $p^E(\theta_1) < p^E(\theta_2)$ for $\theta_1, \theta_2 \in [0, 1]$.

Proof:

a)Firstly, note that if $F_H(v) \leq F_L(v)$, then $F_H(v) \leq \theta F_H(v) + (1-\theta) F_L(v) \leq F_L(v)$ for any $\theta \in [0, 1]$. Secondly, note that $G(v) = N_s F(v)^{N-1} - (N_s - 1) F(v)^N$ is a monotone increasing function of $F(v)$ and, therefore, preserves the relationship $F_H(v) \leq \theta F_H(v) + (1-\theta) F_L(v) \leq F_L(v)$ for any $\theta \in [0, 1]$. Hence, we conclude that $G_H(v) \leq G(v, \theta) \leq G_L(v)$ and $\int_0^V v dG_L(v) \leq \int_0^V v dG(v, \theta) \leq \int_0^V v dG_H(v)$ for any $\theta \in [0, 1]$. This establishes that if $F_H(v) \leq F_L(v)$, then $p_L^P \leq p^E(\theta) \leq p_H^P$ for $\theta \in [0, 1]$.

b) Similarly to the proof of part (a), note that if $\theta_1 < \theta_2$, then $G(v, \theta_1) < G(v, \theta_2)$ by increasing monotonicity of $G(v)$. Hence, we conclude that $p^E(\theta_1) < p^E(\theta_2)$.■

The results in Lemma 3.1 suggests that when the distribution function of valuations for an item of high opaque quality first-order stochastically dominates the distribution function of valuations for an item of low opaque quality, the expected price in an electronic platform is at least as high as the expected price for an item of low opaque quality in a physical platform and at most as high as the expected price for an item of high opaque quality in a physical platform. In part (b) of Lemma 3.1 we show that the expected price in an electronic platform should increase if bidders attach a higher probability to a high opaque quality item. In the next proposition we show that a seller always lists an item of low opaque quality in an electronic platform.

Proposition 3.1

It is a dominant strategy for a seller to list an item of low opaque quality (type q_L) in an electronic platform.

Proof:

Note that the maximum expected payoff a seller can obtain by listing an item of type q_L in a physical auction is $(1-\delta)p_L^P$, where $\delta > 0$ is the size of the commission in a physical platform. Note further that by Lemma 3.1, $(1-\delta)p_L^P < p_L^P < p^E(\theta)$ for $\delta > 0$ and any bidders' belief $\theta \in [0,1]$. Hence, a seller obtains a strictly higher payoff by listing an item of type q_L in an electronic platform for any bidders' belief, and it is a dominant strategy for a seller of a low quality item to list only in an electronic platform. ■

In the next proposition we derive equilibrium conditions for sorting of items of high opaque quality between electronic and physical platforms. We use the notion of a Perfect Bayesian Equilibrium (PBE) to derive results in Proposition 3.2.

Proposition 3.2

Let p_H^P be the expected price for a type q_H item in a physical platform, p_L^P be the expected price for a type q_L item in a physical platform, and $p^E(\theta = \alpha)$ be the expected price for an item in an electronic platform when bidders believe that the probability of facing a type q_H item in an electronic platform is α, where α is determined by nature. Then the following holds:

(a) If $p^E(\theta = \alpha) \geq (1-\delta)p_H^P \geq p_L^P$, then there are two PBE: (1) a seller lists both a type q_H item and a type q_L item in an electronic platform, and (2) a seller lists a type q_H item in a physical platform and a type q_L item in an electronic platform.

(b) If $p^E(\theta = \alpha) \geq p_L^P \geq (1-\delta)p_H^P$, then there is a unique pooling PBE and a seller lists both a type q_H item and a type q_L item in an electronic platform.

(c) If $(1-\delta)p_H^P \geq p^E(\theta = \alpha) \geq p_L^P$, then there is a unique separating PBE and a seller lists a type q_H item in a physical platform and a type q_L item in an electronic platform.

Proof:

Note that by Proposition 3.1, $p^E(\theta = \alpha) \geq p_L^P$ is always true. Hence, we need to establish equilibrium seller's strategies constitute when $p^E(\theta = \alpha) \geq (1-\delta)p_H^P \geq p_L^P$, $p_L^P \geq (1-\delta)p_H^P$, and when $(1-\delta)p_H^P \geq p^E(\theta = \alpha)$.

Before we proceed further, recall that $\beta \in [0,1]$ is the seller's probability of listing a type q_H item in a physical auction and $(1-\beta)$ is the sellers' probability of listing a type q_H item in an electronic auction. Similarly, $\gamma \in [0,1]$ is the seller's probability of listing a type q_L item in a physical auction and $(1-\gamma)$ is the seller's probability of listing a type q_L item in an electronic auction. Given the seller's strategies, we can define bidders' belief θ about a type q_H item listed in an electronic auction. We assume that bidders form their belief θ about a type q_H item in an electronic platform according to Bayes' rule. In particular, we assume that $\theta(\beta,\gamma) = \frac{\alpha(1-\beta)}{\alpha(1-\beta)+(1-\alpha)(1-\gamma)}$, where α is the probability of a type q_H item as determined by nature, $\alpha(1-\beta)$ is the probability that a type q_H item is listed in an electronic auction given the seller's behavioral strategy (β,γ), $\alpha(1-\beta)+(1-\alpha)(1-\gamma)$ is the probability of listing both types of items in an electronic auction given the seller's behavioral strategy (β,γ). Similarly, the probability that an item listed in an electronic auction is of type q_L is $1-\theta(\beta,\gamma) = \frac{(1-\alpha)(1-\gamma)}{\alpha(1-\beta)+(1-\alpha)(1-\gamma)}$.

By Proposition 3.1, a seller always lists a type q_L item in an electronic auction, which implies that $\gamma = 0$ for any belief θ. Given that $\gamma = 0$, we need to consider only three possible cases: (1) $\{\beta = 1, \gamma = 0\}$, (2) $\{\beta = 0, \gamma = 0\}$, and (3) $\{\beta \in (0,1), \gamma = 0\}$. Note that expected prices for a type q_H item and a type q_L item in a physical auction do not depend on seller's behavioral strategies. Hence, in all three cases we denote the expected price in a physical auction for a type q_H item by p_H^P and the expected price in a physical auction for a type q_L item by p_L^P.

Case (1): Given the seller's strategy of $\beta = 1$ and $\gamma = 0$, bidder belief about a type q_H item in an electronic auction is $\theta(\beta = 1, \gamma = 0) = 0$. As a result, the expected price in an electronic auction is $p^E(\theta = 0) = p_L^P$. Given this price, the incentive compatibility constraint for a seller listing a type q_H item is $IC_H : (1-\delta)p_H^P \geq \beta(1-\delta)p_H^P + (1-\beta)p_L^P$, where the left-hand side of the inequality is the seller's payoff from following the behavioral strategy $\beta = 1$ and the right-hand side is the deviation payoff. Similarly, the incentive compatibility constraint for a seller listing a type q_L item is $IC_L : p_L^P \geq \gamma(1-\delta)p_L^P + (1-\gamma)p_L^P$, where the left-hand side is the seller's payoff from following the behavioral strategy $\gamma = 0$ and the right-hand side is the deviation payoff. Note that the incentive compatibility constraint IC_L is always satisfied and the incentive compatibility constraint IC_H is satisfied if $(1-\delta)p_H^P \geq p_L^P$. Hence, the seller's strategy $\{\beta = 1, \gamma = 0\}$ is an equilibrium if $(1-\delta)p_H^P \geq p_L^P$.

Case (2): Given the seller's strategy of $\beta = 0$ and $\gamma = 0$, bidder belief about a type q_H item in an electronic auction is $\theta(\beta = 0, \gamma = 0) = \alpha$. As a result, the expected price in an electronic auction is $p^E(\theta = \alpha)$. Given this price, the incentive compatibility constraint for a seller listing a type q_H item is $IC_H : p^E(\theta = \alpha) \geq \beta(1-\delta)p_H^P + (1-\beta)p^E(\theta = \alpha)$, where the left-hand side of the inequality is the seller's payoff from following the behavioral strategy $\{\beta = 0, \gamma = 0\}$ and the right-hand side is the deviation payoff. Similarly, the incentive compatibility constraint for a seller listing a type q_L item is $IC_L : p^E(\theta = \alpha) \geq \gamma(1-\delta)p_L^P + (1-\gamma)p^E(\theta = \alpha)$, where the left-hand side is the seller's payoff from following the behavioral strategy $\{\beta = 0, \gamma = 0\}$ and the right-hand side is the deviation payoff. Since the right-hand side of IC_L is less than the right-hand side of IC_H, both IC_L and IC_H hold if IC_H holds. Since IC_H holds if $p^E(\theta = \alpha) \geq (1-\delta)p_H^P$, the seller's behavioral strategy $\{\beta = 0, \gamma = 0\}$ is an equilibrium if $p^E(\theta = \alpha) \geq (1-\delta)p_H^P$.

Case (3): Given the seller's behavioral strategy of $\beta \in (0,1)$ and $\gamma = 0$, bidder belief about a type q_H item in an electronic auction is $\theta(\beta \in (0,1), \gamma = 0) = \frac{\alpha(1-\beta)}{\alpha(1-\beta)+(1-\alpha)}$. As a result, the expected price in an electronic auction is $p^E(\theta = \frac{\alpha(1-\beta)}{\alpha(1-\beta)+(1-\alpha)})$. Given this price, the seller's payoff from randomizing with probability $\beta \in (0,1)$ between listing a type q_H item in a physical platform and an electronic platform is $\beta(1-\delta)p_H^P + (1-\beta)p^E(\theta = \frac{\alpha(1-\beta)}{\alpha(1-\beta)+(1-\alpha)})$. The deviation payoff depends on whether $(1-\delta)p_H^P > p^E(\theta = \frac{\alpha(1-\beta)}{\alpha(1-\beta)+(1-\alpha)})$. Let's consider both cases and first assume that $(1-\delta)p_H^P > p^E(\theta = \frac{\alpha(1-\beta)}{\alpha(1-\beta)+(1-\alpha)})$. Then the incentive compatibility constraint for a seller listing a type q_H item is $IC_H : \beta(1-\delta)p_H^P + (1-\beta)p^E(\theta = \frac{\alpha(1-\beta)}{\alpha(1-\beta)+(1-\alpha)}) \geq (1-\delta)p_H^P$. Note that this IC_H constraint is violated, since $(1-\delta)p_H^P > p^E(\theta = \frac{\alpha(1-\beta)}{\alpha(1-\beta)+(1-\alpha)})$.

Let's consider the second case and assume that $(1-\delta)p_H^P < p^E(\theta = \frac{\alpha(1-\beta)}{\alpha(1-\beta)+(1-\alpha)})$. Then the incentive compatibility constraint for a seller listing a type q_H item is $IC_H : \beta(1-\delta)p_H^P + (1-\beta)p^E(\theta = \frac{\alpha(1-\beta)}{\alpha(1-\beta)+(1-\alpha)}) \geq p^E(\theta = \frac{\alpha(1-\beta)}{\alpha(1-\beta)+(1-\alpha)})$. Similarly, this IC_H constraint is violated, since $(1-\delta)p_H^P < p^E(\theta = \frac{\alpha(1-\beta)}{\alpha(1-\beta)+(1-\alpha)})$. Hence, the seller's strategy of $\beta \in (0,1)$ and $\gamma = 0$ is not an equilibrium.

Further, note that equilibria in cases (1) and (2) hold together when $p^E(\theta = \alpha) \geq (1-\delta)p_H^P \geq p_L^P$, the equilibrium in case (1) is unique when $p^E(\theta = \alpha) \geq p_L^P \geq (1-\delta)p_H^P$, and the equilibrium in case (2) is unique when $(1-\delta)p_H^P \geq p^E(\theta = \alpha) \geq p_L^P$. This concludes the proof of Proposition 3.2. ■

The results in Proposition 3.2 suggest that depending on the commonly known distribution of high and low opaque quality items, a seller of a high opaque quality can list her item with certainty either in an physical platform or in an electronic platform. This is in contrast to the result in Proposition 3.1, where we show that a seller of a low opaque quality item always lists her item only in an electronic platform.

When the condition in part (c) of Proposition 3.2 holds, a seller of a high opaque quality item always lists her item in a physical auction, and we have pure market segmentation with multiple identical sellers of high opaque quality items listing their items in a physical platform and multiple identical sellers of low opaque quality items listing their items in an electronic platform. When the condition in part (b) of Proposition 3.2 holds, the physical platform collapses, since multiple identical sellers of both high and low opaque quality items list only in an electronic platform. When the condition in part (a) of Proposition 3.1 holds, we have an impure market segmentation with possibly some identical sellers of high opaque quality items listing their items in an electronic platform and some in a physical platform and all sellers of low opaque quality items listing their items in an electronic platform.

Note that the seller's equilibrium listing strategy does not depend on the total number of bidders in each platform as long as the number of bidders in an electronic and a physical platform is the same. If the number of bidders in electronic and physical platforms is different, then the ranking of prices in Lemma 3.1 may not hold, and the results in Propositions 3.1 and 3.2 may not be valid.

Another important implication of Proposition 3.2 is that whenever a physical platform exists for single-dimensional quality items, the quality of items in a physical platform is no worse than the quality of items listed in an electronic platform. We formally state this implication in Corollary 3.1.

Corollary 3.1.

For single-dimensional quality items, whenever a physical platform exists, the quality of items listed in a physical platform is no worse than the quality of items listed in an electronic platform.

3.5 Two-dimensional Quality

In this section we generalize the model and introduce two-dimensional quality. Consider an item of quality consisting of two parameters: a transparent quality parameter and an opaque quality parameter. For simplicity we assume that the transparent quality parameter, t, can be either high or low, $t \in \{H, L\}$. The opaque quality parameter, i, can also be high or low, $i \in \{H, L\}$. Depending on whether the transparent and the opaque quality parameters are high or low, an item can be one of four possible types: $q_{H,H}$, $q_{H,L}$, $q_{L,H}$, $q_{L,L}$, where the first subscript indicates the transparent quality and the second subscript indicates the opaque quality. As before, we assume that the opaque quality is not observable in an electronic platform but perfectly observable in a physical platform. However, the transparent quality is perfectly observable both in an electronic and in a physical platform. The probability of each type is determined by nature and is common knowledge in the model[5]. We denote the probability of a type $q_{t,i}$ item by $\alpha_{t,i}$, $t, i \in \{H, L\}$, and assume that $\sum\limits_{t,i \in \{H,L\}} a_{t,i} = 1$.

Bidders derive their valuations for each type of item from a corresponding stochastically independent cumulative distribution function. We denote a cumulative distribution function of valuations for an item of type $q_{t,i}$ by $F_{t,i}(v)$, $v \in [0, V]$ and $t, i \in \{H, L\}$. As before, we assume that conditional on transparent quality a cumulative distribution function of valuations for a high opaque quality item first-order stochastically dominates a cumulative distribution function of a low opaque quality item, or that $F_{t,H}(v) \leq F_{t,L}(v)$ for all $v \in [0, V]$ and $t \in \{H, L\}$. In addition, we assume that conditional on opaque quality a cumulative distribution function of valuations for a high transparent quality item first-order stochastically dominates a cumulative distribution function of a low transparent quality item, or that $F_{H,i}(v) \leq F_{L,i}(v)$ for all $v \in [0, V]$ and $i \in \{H, L\}$.

[5] In other words, we assume that bidders know the probability of an item of each type before a seller chooses a specific sale platform.

We define a menu of prices for each type of item in a physical platform and an electronic platform. The menu of prices for a high transparent quality item is $\{p^P_{H,H}, p^P_{H,L}, p^E_H(\theta_H)\}$, where the superscript denotes whether the price is formed in a physical or in an electronic platform, the first subscript letter "H" in all three prices indicates that the prices belong to a high transparent quality item, and the second subscript letters in the first two prices indicate whether the opaque quality is high or low. The argument θ_H in the third price indicates the bidders' belief that a high transparent quality item has high opaque quality. The menu of prices for a low transparent quality item $\{p^P_{L,H}, p^P_{L,L}, p^E_L(\theta_L)\}$ is defined in a similar fashion with the only difference that θ_L indicates the bidders' belief that a low transparent quality item has high opaque quality. By part (a) of Lemma 3.1, we have that $p^P_{H,L} \leq p^E_H(\theta_H) \leq p^P_{H,H}$, $p^P_{L,L} \leq p^E_L(\theta_L) \leq p^P_{L,H}$, $p^P_{L,L} \leq p^P_{L,H} \leq p^P_{H,H}$, and $p^P_{L,L} \leq p^P_{H,L} \leq p^P_{H,H}$.

We can derive sorting conditions of items of different quality between electronic and physical platforms. Since bidders observe the transparent quality of items in both platforms, the market for items with two-dimensional quality essentially breaks into two separate segments: with high transparent quality items and with low transparent quality items. Hence, the results of Proposition 3.1, Proposition 3.2 and Corollary 3.1 about sorting of opaque quality items are true for each segment. Further, the sorting of items across two platforms does not depend on the transparent quality *per se*, since by Proposition 3.1 all low opaque quality items are listed in an electronic platform, and by Proposition 3.2 the sorting of high opaque quality items only depends on the probabilities $\alpha_{H,H}$, $\alpha_{H,L}$, $\alpha_{L,H}$, and $\alpha_{L,L}$.

In the next two corollaries we derive conclusions about the quality of items listed in the two platforms. In Corollary 3.2 we state that the opaque quality of high transparent quality items in a physical platform is no worse than the opaque quality of low transparent quality items in a physical platform, if the physical platform exists. This is a direct consequence of Proposition 3.1: Since only low opaque quality items are always listed in an electronic platform, in a physical platform the opaque quality of high transparent quality items is no worse than the opaque quality of low transparent quality items.

Without making any additional assumptions about the number of items of each type and quality indexes, in Corollary 3.3 we present sufficient conditions when the opaque quality of high transparent quality items in an electronic platform is no worse than the opaque quality of low transparent quality items in an electronic platform, and when both the opaque and the transparent quality of items in a physical platform is no worse than the opaque and the transparent quality of items in an electronic platform.

Corollary 3.2.

If a physical platform exists for items of high and low transparent quality, then the opaque quality of high transparent quality items listed in a physical platform is no worse than the opaque quality of low transparent quality items listed in a physical platform.

Corollary 3.3.

Let $\alpha_{t,i}$ be the probability of an item of type $q_{t,i}$ and $p_t^{E[-1]}(.)$ denote an inverse of a price in an electronic platform, $t, i \in \{H, L\}$.

a) The opaque quality of high transparent quality items listed in an electronic platform is no worse than the opaque quality of low transparent quality items listed in an electronic platform if $\frac{\alpha_{L,H}}{\alpha_{L,H}+\alpha_{L,L}} < \frac{\alpha^*_{L,H}}{\alpha^*_{L,H}+\alpha^*_{L,L}} = p_L^{E[-1]}((1-\delta)p^P_{L,H})$.

b) The opaque and the transparent quality of items listed in a physical platform is no worse than the opaque and the transparent quality of items listed in an electronic platform if $\frac{\alpha_{H,H}}{\alpha_{H,H}+\alpha_{H,L}} < \frac{\alpha^*_{H,H}}{\alpha^*_{H,H}+\alpha^*_{H,L}} = p_H^{E[-1]}((1-\delta)p^P_{HH})$ and $p^P_{L,L} \geq (1-\delta)p^P_{L,H}$.

Proof:

a) By Proposition 3.1, low opaque quality items of any transparent quality are exclusively listed in an electronic platform. Next, consider the equilibrium listing strategies of sellers of high opaque quality items of any transparent quality. Consider a pair of probabilities ($\alpha^*_{L,H}$, $\alpha^*_{L,L}$) for a low transparent quality item such that $(1-\delta)p^P_{L,H} = p^E_L(\theta_L =$

$\frac{\alpha^*_{L,H}}{\alpha^*_{L,H}+\alpha^*_{L,L}}$). Then we can define an inverse $\frac{\alpha^*_{L,H}}{\alpha^*_{L,H}+\alpha^*_{L,L}} = p_L^{E[-1]}((1-\delta)p^P_{L,H})$ by continuity of $p^E_L(\theta_L)$ for any $F(v) > 0$ and $N \geq 0$. Since by part (b) of Lemma 3.1, $p^E_L(\theta_L)$ is monotone increasing in θ_L, by part (c) of Proposition 3.2 we have that for any pair of probabilities $(\alpha_{L,H}, \alpha_{L,L})$, such that $\frac{\alpha_{L,H}}{\alpha_{L,H}+\alpha_{L,L}} < \frac{\alpha^*_{L,H}}{\alpha^*_{L,H}+\alpha^*_{L,L}}$, sellers always list high opaque and low transparent quality items in a physical platform. This means that the only low transparent quality items listed in an electronic platform are the ones with low opaque quality. Hence, for any listing strategy of sellers of high transparent quality items, the opaque quality of high transparent quality items listed in an electronic platform is no worse than the opaque quality of low transparent quality items listed an electronic platform.

b) Similarly to the proof of part (a), note that for any pair of probabilities $(\alpha_{H,H}, \alpha_{H,L})$, such that $\frac{\alpha_{H,H}}{\alpha_{H,H}+\alpha_{H,L}} < \frac{\alpha^*_{H,H}}{\alpha^*_{H,H}+\alpha^*_{H,L}} = p_H^{E[-1]}((1-\delta)p^P_{HH})$, sellers always list high opaque and high transparent quality items in a physical platform. By part (c) of Proposition 3.2, when $p^P_{L,L} \geq (1-\delta)p^P_{L,H}$, sellers always list high opaque and low transparent quality items in an electronic platform. Given that by Proposition 3.1, low opaque quality items of any transparent quality are always listed in an electronic platform, the quality of items in a physical platform is no worse than the quality of items in an electronic platform. ■

Part (a) of Corollary 3.3 states that unless all high opaque and low transparent quality items are listed in a physical platform, we cannot guarantee that the opaque quality of high transparent quality items in an electronic platform is higher than the opaque quality of low transparent quality items in an electronic platform. This result means that there is necessarily a positive correlation between transparent and opaque quality in an electronic platform only if all low transparent and high opaque items are listed in a physical platform, which can happen only if bidders have a sufficiently low belief that a low transparent quality item has high opaque quality.

Part (b) of Corollary 3.3 states that the opaque and the transparent quality of items in a physical platform is no worse than the opaque and the transparent quality of items in an electronic platform, if items of high transparent and high opaque quality are exclusively listed in a physical platform and if items of low transparent and high opaque quality are exclusively listed in an electronic platform. In other words, the overall quality of items, as defined as opaque and transparent quality, is necessarily higher in a physical platform than in an electronic if the conditions in Part (b) of Corollary 3.3 are satisfied. More specific conditions on the quality sorting of items of different types in electronic and physical platforms require assumptions on the number of items of each type and the quality indexes.

The main conclusion in this section is that the introduction of an additional quality dimension breaks down the unambiguous quality sorting implication of the basic model with one-dimensional quality. Unless restrictive conditions of Corollary 3.3 are satisfied, without any additional assumptions it is impossible to make any conclusions about the quality sorting across different platforms conditionally on some specific quality parameter or unconditionally on any quality parameters. To illustrate this point, consider a plausible scenario when items of low transparent and high opaque quality are listed in a physical platform, while items of high transparent and high opaque quality together with items of high transparent and low opaque quality and items of low transparent and low opaque quality are listed in an electronic platform. In this scenario, it is impossible to make any conclusion about the average quality of items in both platforms without any further information about the quantity of items of each type and the quality values.

3.6 Discussion

We can draw several general empirical implications from our model, which should hold independently of bidders' beliefs or the dimensionality of quality. First of all, by Corollary 3.1, holding the transparent quality constant, the opaque quality of items in a physical platform should be no worse than the opaque quality of items in an electronic platform.

Secondly, by Corollary 3.2, we should observe a non-negative correlation between transparent and opaque quality for items in a physical platform. Lastly, as Corollary 3.3 shows, it is impossible to establish the correlation between transparent and opaque quality for items in an electronic platform without an additional structure and without some restrictions on bidders' beliefs. For the same reason, it is impossible to conclude how the overall (opaque and transparent) quality of items in a physical platform compares to the overall quality of items in an electronic platform. However, if we observe a positive correlation between transparent and opaque quality in an electronic platfom, it should be the case that the share of items with low transparent and high opaque quality in an electronic platform is either low or non-existent. By Proposition 3.2, we know that low transparent and high opaque quality items are never listed in an electronic platform if the bidders' belief about high opaque quality of a low transparent quality item is sufficiently low. Similarly, if we observe a negative correlation between transparent and opaque quality in an electronic platfom, it should be the case that the share of items with low transparent and high opaque quality in an electronic platform is sufficintly high. By Proposition 3.2, we know that this can happen only if low transparent and high opaque quality items are listed in an electronic platform, which is possible if the bidders' belief about high opaque quality of a low transparent quality item is sufficiently high.

The closest available study of the quality sorting between electronic and physical platforms is the work by Jin and Kato (2007). The authors study the sorting of graded and ungraded baseball cards of different quality between online market and a retail market. The authors consider a model with one-dimensional quality, grading costs, retail listing fees, and unobservable quality of ungraded cards in online market[6]. The authors derive the quality ranking of cards of different quality across different platforms and find that the quality of graded cards traded online is no worse than the quality of ungraded cards traded

[6]The authors assume that the search costs in online market are zero, while the search costs in a retail market are a fixed share of the sale price. Hence, the search costs in a retail market essentially act like the commission in our physical platform.

in a retail market, which in turn is no worse than the quality of ungraded cards traded online. This result is consistent with the conclusions of our model with one-dimensional quality.

In our model, we would combine graded cards sold online and the cards sold in a retail market into one category, because the quality of graded cards listed online and the quality of cards sold in a retail market is perfectly observable and because there are costs associated with grading and with listing in a retail market. Since our model suggests that the lowest quality baseball cards will always be listed online ungraded, the quality of cards listed online graded and the quality of cards listed in a retail market ungraded will always be higher. This is consistent with results in Jin and Kato (2007). However, as we show in the rest of the chapter, when we introduce an additional quality dimension, the clear quality sorting result in Jin and Kato (2007) and in our one-dimensional quality model may not hold.

We consider an extension of the two-dimensional quality model and introduce quality grading. An introduction of the quality grading in the two-dimensional quality model adds one additional platform to the existing electronic and physical platforms. Hence, high and low transparent quality items of high and low opaque quality can be listed in a physical platform, ungraded in an electronic platform, and graded in an electronic platform[7]. Since grading essentially reveals the quality of an item online, a seller will grade an item and list it in an electronic platform if the grading cost is less than the physical platform listing fee. Otherwise, a seller will list her item in a physical platform ungraded. As a result, whenever grading is available, an electronic platform with graded items will replace a physical platform if the physical platform listing fee is more than the grading cost, and a physical platform will replace an electronic platform with graded items if the grading cost is more than the physical platform listing fee.

[7] A seller will not grade an item listed in a physical platform, because grading is costly and because the quality of an item listed in a physical platform is assumed to be perfectly observable with or without grading.

The coexistence of a physical platform and an electronic platform with graded items is possible if we introduce an additional intermediate opaque quality with a property that the expected payoff from listing an ungraded intermediate quality item in a physical platform is higher than the expected payoff from grading and listing the intermediate quality item in an electronic market. In addition, it must be the case that the expected payoff from listing a graded high quality item in an electronic platform should exceed the expected payoff from listing the high quality item ungraded in a physical platform. This modification is consistent with the model of Jin and Kato (2007), who assume a continuous one-dimensional quality. Note that the presence of a transparent quality dimension in our model with grading does not have any role in determining listing patterns. Another way through which we can introduce the coexistence of a physical platform and an electronic platform with graded items is by introducing the mechanism of Ellison, Fudenberg and Mobius (2004) and allow sellers to affect seller-buyer platform ratios by choosing one platform over another.

The second possible extension is the introduction of heterogeneity in bidders' distribution functions of valuations. Note that the introduction of heterogeneity in bidders' distribution functions does not have any impact on our equilibrium predictions as long as the ranking of prices in Lemma 3.1 is preserved. The ranking of prices in Lemma 3.1 solely relies on the assumption of the first-order stochastic dominance. Hence, as long as the distribution of the second-highest order statistic for the high opaque quality item first-order stochastically dominates the distribution of the second-highest order statistic for the low opaque quality item, all results in our model are preserved independently of the degree of heterogeneity in bidders' distribution functions.

3.7 Concluding Remarks

In this chapter we construct a simple model of asymmetric information, where a seller of an item with unobservable quality parameter chooses between electronic and physical platforms. The physical platform is costly to use and completely discloses the quality parameter. The electronic platform is free to use and does not reveal the quality parameter. A seller sends an informative signal about quality of her item to potential buyers by choosing one platform over another. We use Spence's (1973) signalling model as a general framework for the analysis. We find that a seller always lists low unobservable quality item in an electronic platform. However, a higher unobservable quality item can be listed both in electronic and in a physical platform. As a result, we conclude that the quality of items in a physical platform should be no worse than the quality of items in an electronic platform. However, if we introduce an additional transparent quality parameter, the quality ranking between the two platforms no longer holds.

References

Aghion, Philippe, and Patrick Bolton. "Contracts as a Barrier to Entry." *American Economic Review* 77 (1987): 388-401.

Aghion, Philippe, Mathias Dewatripont, and Patrick Rey. "Renegotiation Design with Unverifiable Information." *Econometrica* 62 (1994): 257-282.

Akerlof, George A. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84 (1970): 488–500.

Anderson, Simon P., and Stephen Coate. "Market Provision of Broadcasting: A Welfare Analysis." *Review of Economic Studies* 72 (2005): 947–72.

Aoki, Masahiko. "Toward an Economic Model of the Japanese Firm." *Journal of Economic Literature* 28 (1990): 1-27.

Armstrong, Mark. "Competition in Two-Sided Markets." *RAND Journal of Economics* 37 (2006): 668–91.

Bajari, Patrick. and Ali Hortacsu. "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions." *RAND Journal of Economics* 34 (2003): 329–355.

Bergemann, Dirk and Said Maher. "Dynamic Auctions." <u>Wiley Encyclopedia of Operations Research and Management Science.</u> New York, NY: Wiley, 2011.

Binmore, Kenneth, Avner Shaked, and John Sutton. "An Outside Option Experiment." *Quarterly Journal of Economics* 104 (1989): 753-770.

Cai, H., Riley, J., and Ye, L. "Reserve Price Signaling." *Journal of Economic Theory* 135 (2007): 253-268.

Carare, O. "Reserve Prices in Repeated Auctions." *Review of Industrial Organization* 40 (2012): 225–247.

Casadesus-Masanell, Ramon, and Daniel Spulber. "The Fable of Fisher Body." *Journal of Law and Economics* 43 (2000): 67-104.

Che, Yeon-Koo, and Donald Hausch. "Cooperative Investments and the Value of Contracting." *American Economic Review* 89 (1999): 125–147.

Che, Yeon-Koo, and József Sákovics. "A Dynamic Theory of Holdup." *Econometrica* 72 (2004): 1063-1103.

Coase, Ronald. "The Nature of the Firm." *Economica* 4 (1937): 386-405.

Coase, Ronald. "The Acquisition of Fisher Body by General Motors." *Journal of Law and Economics* 43 (2000): 15-31.

Edlin, Aaron, and Stefan Reichelstein. "Holdups, Standard Breach Remedies, and Optimal Investment." *American Economic Review* 86 (1996): 478-501.

Einav, E., Kuchler, T., Levin, J. and Sundaresan, N. "Learning from Seller Experiments in Online Markets." working paper (2012).

Ellison, Glenn, and Drew Fudenberg. "Knife-Edge or Plateau: When Do Market Models Tip?" *Quarterly Journal of Economics* 118 (2003): 1249–78.

Ellison, Glenn, Drew Fudenberg, and Markus Mobius. "Competing Auctions." *Journal of the European Economic Association* 2 (2004): 30–66.

Freeland, Robert. "Creating Holdup Through Vertical Integration: Fisher Body Revisited." *Journal of Law and Economics* 43 (2000): 33-66.

Grout, Paul. "Investment and Wages in the Absence of Binding Contract: A Nash Bargaining Approach." *Econometrica* 52 (1984): 449-460.

Grossman, Sanford, and Oliver Hart. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Investment." *Journal of Political Economy* 94 (1986): 691-719.

Guriev, Sergei, and Dmitriy Kvasov. "Contracting on Time." *American Economic Review* 95 (2005): 1369-1385.

Hart, Oliver, and John Moore. "Incomplete Contracts and Renegotiation." *Econometrica* 56 (1988): 755-85.

Hasker, K. and Sickles, R. "eBay in the Economic Literature: Analysis of an Auction Marketplace." *Review of Industrial Organization* 37 (2010): 3–42.

Hidvegi, Z., Wang, W., and Whinston, A. "Buy-Price English Auction." *Journal of Economic Theory* 129 (2006): 31-56.

Holmström, Bengt, and John Roberts. "The Boundaries of the Firm Revisited." *Journal of Economic Perspectives* 12 (1998): 73-94.

Jin, Ginger Z., and Andrew Kato. "Dividing Online and Offline: A Case Study." *Review of Economic Studies* 74 (2007): 981-1004.

Joskow, Paul. "Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets." *American Economic Review* 77 (1987): 168-185.

Kagel, John. and Daniel Levin. Common Value Auctions and the Winner's Curse. Princeton, NJ: Princeton University Press, 2002.

Katkar, Rama, and David Reiley. "Public Versus Secret Reserve Prices in eBay Auctions: Results from a Pokemon Field Experiment." *Advances in Economic Analysis and Policy* 6 (2006): 1-23.

Klein, Benjamin, Robert Crawford, and Armen Alchian. "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process." *Journal of Law and Economics* 21 (1978): 297-326.

Klein, Benjamin. "Fisher-General Motors and the Nature of the Firm." *Journal of Law and Economics* 43 (2000): 105-141.

Krishna, Vijay. Auction Theory, 2nd. edition. Academic Press, 2010.

Lafontaine, Francine, and Margaret Slade. "Vertical Integration and Firm Boundaries: The Evidence." *Journal of Economic Literature* 45 (2007): 629-685.

Levin, D. and Smith, J. "Optimal Reservation Prices in Auctions." *The Economic Journal* 106 (1996): 1271-1283.

Lewis, Gregory. "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors." *American Economic Review* 101 (2011): 1535-46.

Lucking-Reiley, D., Bryan, D., Prasad, N., and Reeves, D. "Pennies from eBay: the Determinants of Price in Online Auctions." *Journal of Industrial Economics* 55 (2007): 223–233.

Maskin, Eric, and John Riley. "Monopoly with Incomplete Information." *RAND Journal of Economics* 15 (1984): 171–96.

Mathews, T., "A Risk Averse Seller in a Continuous Time Auction with a Buyout Option." *Brazilian Electronic Journal of Economics* 5 (2003): 1-26.

McAffee, P. and Vincent, D. "Sequentially Optimal Auctions." *Games and Economic Behavior* 18 (1997): 246-276.

Milgrom, Paul and Robert Weber. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (1982): 1089–1122.

Mussa, Michael, and Sherwin Rosen. "Monopoly and Product Quality." *Journal of Economic Theory* 18 (1978): 301–17.

Myerson, Roger B. "Optimal Auction Design." *Mathematics of Operations Research* 6 (1981): 58–73.

Nöldeke, Georg and Klaus Schmidt. "Option Contracts and Renegotiations: A Solution to the Hold-up Problem." *RAND Journal of Economics* 26 (1995): 163 - 179.

Ockenfels, A. and Roth, A. "Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction." *Games and Economic Behavior* 55 (2006): 297–320.

Osborne, Martin, and Ariel Rubinstein. Bargaining and Markets. San Diego, CA: Academic Press, Inc., 1990.

Overby, Eric M. and Jap, Sandy D. "Electronic and Physical Market Channels: A Multi-Year Investigation in a Market for Products of Uncertain Quality." *Management Science* 55 (2009): 940-957.

Riley, J., and Samuelson,W. "Optimal Auctions." *American Economic Review* 71 (1981): 381–392.

Rochet, Jean-Charles, and Jean Tirole. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association* 1 (2003): 990–1029.

Rosenkranz, S., and Schmitz, P. W. “Reserve Prices in Auctions as Reference Points.” *The Economic Journal* 117 (2007): 637–653.

Rubinstein, Ariel. "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50 (1982): 97-109.

Sloof, Randolph, Joep Sonnemans, and Hessel Oosterbeek. "Specific Investments, Holdup, and the Outside Option Principle." *European Economic Review* 48 (2004): 1399 – 1410.

Spence, Michael. "Job Market Signaling." *Quarterly Journal of Economics* 87(1973): 355–374.

Vincent, D. “Bidding Off the Wall: Why Reserve Prices May be Kept Secret.” *Journal of Economic Theory* 65 (1995): 575–584.

Williamson, Oliver. The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting. New York, NY: Free Press, 1985.

Wolf, James R., and Waleed A. Muhanna. “Adverse Selection and Reputation Systems in Online Auctions: Evidence from eBay Motors.” *Proceedings of the 26th International Conference on Information Systems*, Las Vegas, NV (2005): 847-858.

Appendix A

Proofs and Additional Tests

A.1 Appendix for Chapter 1

Proof of proposition 1.1.

For the existence and uniqueness proofs of part 1 see pages 46-48 of Osborne and Rubinstein (1990).

Proof of part 2:

According to Rubinstein (1982), a pair of SPNE strategies is the solution $(\overline{x}, \overline{y})$ to the system: $\overline{y}_1 = v_1(\overline{x}_1, 1)$ and $\overline{x}_2 = v_2(\overline{y}_2, 1)$, where $v_1(.,1)$ and $v_2(.,1)$ are present value functions of the payoffs obtained with one period of delay. Consider the seller's offer of the form $\overline{x} = (\alpha(w(\sigma) - \underline{v}), \underline{v} + (1-\alpha)(w(\sigma) - \underline{v}))$ and the buyer's offer of the form $\overline{y} = ((1-\beta)(w(\sigma) - \underline{v}), \underline{v} + \beta(w(\sigma) - \underline{v}))$, where α is the seller's share of the trade surplus above the buyer's outside option and β is the buyer's share of the trade surplus above the buyer's outside option. Given that δ_1 is the depreciation rate on the trade surplus and δ_2 is the depreciation rate on the outside option, the SPNE conditions require that $\underline{v} + (1-\alpha)(w(\sigma) - \underline{v}) = \delta_2 \underline{v} + \delta_1 \beta(w(\sigma) - \underline{v})$ and $(1-\beta)(w(\sigma) - \underline{v}) = \delta_1 \alpha(w(\sigma) - \underline{v})$. Solving for α and β, we obtain $\alpha = \frac{1}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{(w(\sigma)-\underline{v})(1-\delta_1^2)}$ and $\beta = \frac{1}{1+\delta_1} - \frac{\delta_1 \underline{v}(1-\delta_2)}{(w(\sigma)-\underline{v})(1-\delta_1^2)}$. By plugging in α and β into $\overline{x}$ and $\overline{y}$ we obtain $\overline{x} = (\frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}, \frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\underline{v}(\delta_2-\delta_1^2)}{1-\delta_1^2})$ and $\overline{y} = (\frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\delta_1 \underline{v}(1-\delta_2)}{1-\delta_1^2}, \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_1-\delta_1^2+\delta_1\delta_2)}{1-\delta_1^2})$.

Consider the seller's offer $(\frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}, \frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\underline{v}(\delta_2-\delta_1^2)}{1-\delta_1^2})$ and note that when $\underline{v} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})$, $\delta_1 \in (\underline{\delta}_1, \overline{\delta}_1)$, and $\delta_2 \in (\underline{\delta}_2, 1)$, the seller's offer $(\frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}, \frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\underline{v}(\delta_2-\delta_1^2)}{1-\delta_1^2})$ is always accepted by the buyer, because the buyer's share of $\frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\underline{v}(\delta_2-\delta_1^2)}{1-\delta_1^2}$

strictly exceeds the buyer's outside option $\underline{v}$. In all other cases, the buyer rejects the seller's offer $(\frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}, \frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\underline{v}(\delta_2-\delta_1^2)}{1-\delta_1^2})$ and takes the outside option $\underline{v}$. Hence, in all cases except when $\underline{v} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})$, $\delta_1 \in (\underline{\delta}_1, \overline{\delta}_1)$, and $\delta_2 \in (\underline{\delta}_2, 1)$ the seller always proposes the offer $(w(\sigma)-\underline{v}, \underline{v})$,which the buyer does not want to reject in favor of the outside option, and the seller's offer $x = (w(\sigma)-\underline{v}, \underline{v})$ together with the buyer's offer $y = (\delta_1(w(\sigma)-\underline{v}), w(\sigma)-\delta_1(w(\sigma)-\underline{v}))$ constitute a pair of SPNE strategies. To see why these strategies are SPNE, note that if the buyer rejects the seller's offer $x = (w(\sigma)-\underline{v}, \underline{v})$, the highest offer the buyer can get is $w(\sigma)-\delta_1(w(\sigma)-\underline{v})$ with one period of delay. Since $\delta_1[w(\sigma)-\delta_1(w(\sigma)-\underline{v})] < \underline{v}$, whenever $\frac{\delta_1 w(\sigma)}{1+\delta_1} < \underline{v}$, the buyer does not have any incentives to reject the seller's offer and negotiate. To see that the seller's strategy is subgame perfect, consider the seller's incentive to reject the buyer's offer $y = (\delta_1(w(\sigma)-\underline{v}), w(\sigma)-\delta_1(w(\sigma)-\underline{v}))$. In case of the rejection, the seller at most can obtain $w(\sigma)-\underline{v}$ with one period of delay or $\delta_1(w(\sigma)-\underline{v})$. Hence, the seller's threat of rejection of any offer $y_1 < \delta_1(w(\sigma)-\underline{v})$ is credible and the seller's startegy is subgame perfect. To see that these startegies are unique, see pages 767-768 of Binmore et. al. (1989).

Consider the cutoff values $\underline{\delta}_1, \overline{\delta}_1 = \frac{1}{2} \mp \sqrt{\frac{1}{4} - \frac{(1-\delta_2)\underline{v}}{w(\sigma)-\underline{v}}}$. We claim that $\underline{\delta}_1$ and $\overline{\delta}_1$ belong to the interval $(0,1)$. First, note that the term $\frac{1}{4} - \frac{(1-\delta_2)\underline{v}}{w(\sigma)-\underline{v}}$ is always less than $\frac{1}{4}$, which implies that $0 < \underline{\delta}_1, \overline{\delta}_1 < 1$. Next, we show that $\frac{1}{4} - \frac{(1-\delta_2)\underline{v}}{w(\sigma)-\underline{v}} \geq 0$. To see this, note that $\frac{\partial(\frac{1}{4} - \frac{(1-\delta_2)\underline{v}}{w(\sigma)-\underline{v}})}{\partial \underline{v}} < 0$ and when $\underline{v} = \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2}$ (the maximum of $\underline{v}$), the term $\frac{1}{4} - \frac{(1-\delta_2)\underline{v}}{w(\sigma)-\underline{v}} = \delta_1^2 - \delta_1 + \frac{1}{4} \geq 0$ for all $\delta_1 \in (0,1)$.

Consider the cutoff value $\underline{\delta}_2 = [1-\delta_1(1-\delta_1)(\frac{w(\sigma)}{\underline{v}}-1)]$. We claim that $\underline{\delta}_2$ belongs to the interval $(0,1)$ when $\underline{v} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, w(\sigma))$. First, note that $\frac{\partial \underline{\delta}_2}{\partial \underline{v}} = \delta_1(1-\delta_1)\frac{w(\sigma)}{\underline{v}^2} > 0$ ($\underline{\delta}_2$ is strictly increasing in $\underline{v}$). When $\frac{\delta_1}{1+\delta_1}w(\sigma) = \underline{v}$, $\delta_1(1-\delta_1)(\frac{w(\sigma)}{\underline{v}}-1) = \delta_1$, and the threshold value is $\underline{\delta}_2 = 1-\delta_1 > 0$. When $\underline{v} \to w(\sigma)$, $\underline{\delta}_2 = 1-\delta_1(1-\delta_1)(\frac{w(\sigma)}{\underline{v}}-1) \to 1$ and for any $\underline{v} < w(\sigma)$, $\underline{\delta}_2 = 1-\delta_1(1-\delta_1)(\frac{w(\sigma)}{\underline{v}}-1) < 1$. Hence, for any $\underline{v} \in (\frac{\delta_1 w(\sigma)}{1+\delta_1}, w(\sigma))$, we have $\underline{\delta}_2 \in (0,1)$.

Next, we show that $\frac{\delta_1 w(\sigma)}{1+\delta_1} < \frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2}$ for any $\delta_2 > \underline{\delta}_2$. To see this note that $\frac{\delta_1}{1+\delta_1} =$

$\frac{\delta_1-\delta_1^2}{1-\delta_1^2} < \frac{\delta_1-\delta_1^2}{1-\delta_2+\delta_1-\delta_1^2}$ for any $\delta_2 > \delta_1$. Since $\underline{\delta}_2 > \delta_1$ whenever $\underline{v} > \frac{\delta_1}{1+\delta_1}w(\sigma)$, we conclude that when $\delta_2 > \underline{\delta}_2$, $\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2} > \frac{\delta_1 w(\sigma)}{1+\delta_1}$.

To prove the uniqueness of the buyer's and the seller's SPNE strategies, we follow Rubinstein (1982). Define M_s as the supremum and m_s as the infimum of the seller's payoffs over all SPNE payoffs in a game where the seller proposes first. Similarly, define $M_b + \underline{v}$ as the supremum and $m_b + \underline{v}$ as the infimum of the buyer's payoffs over all SPNE payoffs in a game where the buyer proposes first. Next, we show that the following inequalities must hold:

(1) $m_b + \underline{v} \geq w(\sigma) - \delta_1 M_s$

(2) $M_s \leq w(\sigma) - (\delta_1 m_b + \delta_2 \underline{v})$

(3) $m_s \geq w(\sigma) - (\delta_1 M_b + \delta_2 \underline{v})$

(4) $M_b + \underline{v} \leq w(\sigma) - \delta_1 m_s$

If the seller rejects the buyer's offer, then the highest payoff the seller can obtain in the present value terms is $\delta_1 M_s$. This implies that the buyer can obtain at least $w(\sigma) - \delta_1 M_s$, which is the right-hand side of inequality (1). Since $m_b + \underline{v}$ is the infimum of the buyer's SPNE payoffs, inequality (1) should hold.

If the buyer rejects the seller's offer, then the smallest payoff the buyer can obtain in the present value terms is $\delta_1 m_b + \delta_2 \underline{v}$. As a result, at most the seller can obtain is $w(\sigma) - (\delta_1 m_b + \delta_2 \underline{v})$, and inequality (2) should hold.

Similarly, if the buyer rejects the seller's offer, then the highest payoff the buyer can obtain in the present value term is $\delta_1 M_b + \delta_2 \underline{v}$. As a result, the seller can at least obtain $w(\sigma) - (\delta_1 M_b + \delta_2 \underline{v})$ implying that inequality (3) should hold.

If the seller rejects the buyer's offer, then the lowest payoff the seller can obtain in the present value term is $\delta_1 m_s$. This implies that the buyer can obtain at most $w(\sigma) - \delta_1 m_s$, which is the right-hand side of inequality (4). Since $M_b + \underline{v}$ is the supremum of the buyer's SPNE payoffs, inequality (4) should hold.

Since the buyer's and the seller's strategies $x = (\frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}, \frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\underline{v}(\delta_2-\delta_1^2)}{1-\delta_1^2})$ and $y = (\frac{\delta_1(w(\sigma)-\underline{v})}{1+\delta_1} + \frac{\delta_1\underline{v}(1-\delta_2)}{1-\delta_1^2}, \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_1-\delta_1^2+\delta_1\delta_2)}{1-\delta_1^2})$ are subgame perfect, we must have that when $\frac{\delta_1}{1+\delta_1}w(\sigma) < \underline{v} < \frac{\delta_1-\delta_1^2}{1-\delta_2+\delta_1-\delta_1^2}w(\sigma)$, then $m_s \leq \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2} \leq M_s$ and $m_b + \underline{v} \leq \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_1-\delta_1^2+\delta_1\delta_2)}{1-\delta_1^2} \leq M_b + \underline{v}$.

Next, we claim that if $\frac{\delta_1}{1+\delta_1}w(\sigma) < \underline{v} < \frac{\delta_1-\delta_1^2}{1-\delta_2+\delta_1-\delta_1^2}w(\sigma)$, then $m_s = \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2} = M_s$ and $m_b + \underline{v} = \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_1-\delta_1^2+\delta_1\delta_2)}{1-\delta_1^2} = M_b + \underline{v}$. To see this, note that inequalities (2) and (1) together imply that $w(\sigma) - M_s \geq \delta_1 m_b + \delta_2\underline{v} \geq \delta_1(w(\sigma) - \delta_1 M_s) + \underline{v}(\delta_2 - \delta_1)$ or that $\frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2} \geq M_s$. The resulting inequality together with the constraint $\frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2} \leq M_s$ implies that $M_s = \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}$. If we plug $M_s = \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}$ into inequality (1), we obtain $m_b \geq \frac{w(\sigma)-\underline{v}}{1+\delta_1} - \frac{\delta_1\underline{v}(1-\delta_2)}{1-\delta_1^2}$. This inequality together with the constraint $m_b + \underline{v} \leq \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_1-\delta_1^2+\delta_1\delta_2)}{1-\delta_1^2}$ implies that $m_b + \underline{v} = \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_1-\delta_1^2+\delta_1\delta_2)}{1-\delta_1^2}$.

Inequalities (3) and (4) together imply that $m_s \geq w(\sigma) - (\delta_1 M_b + \delta_2\underline{v}) \geq w(\sigma) - \delta_1(w(\sigma) - \delta_1 m_s - \underline{v}) - \delta_2\underline{v}$ or that $m_s \geq \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}$. This inequality together with the constraint $m_s \leq \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}$ implies that $m_s = \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}$. If we plug $m_s = \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_2)}{1-\delta_1^2}$ into inequality (4), we obtain $M_b + \underline{v} \leq \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_1-\delta_1^2+\delta_1\delta_2)}{1-\delta_1^2}$. This inequality together with the constraint $\frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_1-\delta_1^2+\delta_1\delta_2)}{1-\delta_1^2} \leq M_b + \underline{v}$ implies that $M_b + \underline{v} = \frac{w(\sigma)-\underline{v}}{1+\delta_1} + \frac{\underline{v}(1-\delta_1-\delta_1^2+\delta_1\delta_2)}{1-\delta_1^2}$.

Proof of part 3:

Given the seller's participation constraint, the buyer at most can obtain $w(\sigma)$ in the relationship. When $\underline{v} > w(\sigma)$, the buyer is strictly worse off accepting the seller's proposal. Hence, the seller's optimal strategy is to offer any division to the seller and the buyer's optimal strategy is to reject the proposal immediately and take the outside option.

Derivation of $\frac{\partial E[\underline{V}|\underline{V}\in(w(\sigma),V]]}{\partial\sigma}$.

Note that $\frac{\partial E[\underline{V}|\underline{V}\in(w(\sigma),V]]}{\partial\sigma} = \frac{\partial}{\partial\sigma}\frac{\int_{w(\sigma)}^{V}\underline{V}dF_{\underline{V}}(\underline{V})}{1-F_{\underline{V}}(w(\sigma))} = \frac{\partial}{\partial\sigma}\frac{\int_{w(\sigma)}^{V}\underline{V}dF_{\underline{V}}(\underline{V})}{\lambda_4(\sigma)}$. By applying Leibniz integration rule, we obtain

$$\frac{\partial}{\partial\sigma}\frac{\int_{w(\sigma)}^{V}\underline{V}dF_{\underline{V}}(\underline{V})}{\lambda_4(\sigma)}=\frac{-w'(\sigma)w(\sigma)f_{\underline{V}}(w(\sigma))}{\lambda_4(\sigma)}-\frac{\lambda_4'(\sigma)\int_{w(\sigma)}^{V}\underline{V}dF_{\underline{V}}(\underline{V})}{(\lambda_4(\sigma))^2}=$$
$$-\lambda_4'(\sigma)\frac{E[\underline{V}|\underline{V}\in(w(\sigma),V]]-w(\sigma)}{\lambda_4(\sigma)}, \text{ since } \lambda_4'(\sigma)=-f_{\underline{V}}(w(\sigma))w'(\sigma).$$

Simplification of equations (1.6a) and (1.6b).

The first step is to expand the terms $\frac{\partial E[\underline{V}|\underline{V}\in(\frac{\delta_1 w(\sigma^a_{eq})}{1+\delta_1},w(\sigma^a_{eq}))]}{\partial\sigma^a_{eq}}$, $\frac{\partial E[\underline{V}|\underline{V}\in(\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2},w(\sigma^b_{eq}))]}{\partial\sigma^b_{eq}}$, and $\frac{\partial E[\underline{V}|\underline{V}\in(\frac{\delta_1 w(\sigma^b_{eq})}{1+\delta_1},\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2})]}{\partial\sigma^b_{eq}}$.

$\frac{\partial E[\underline{V}|\underline{V}\in(\frac{\delta_1 w(\sigma^a_{eq})}{1+\delta_1},w(\sigma^a_{eq}))]}{\partial\sigma^a_{eq}}=\frac{\partial}{\partial\sigma^a_{eq}}\frac{\int_{\frac{\delta_1 w(\sigma^a_{eq})}{1+\delta_1}}^{w(\sigma^a_{eq})}\underline{V}dF_{\underline{V}}(\underline{V})}{F_{\underline{V}}(w(\sigma^a_{eq}))-F_{\underline{V}}(\frac{\delta_1 w(\sigma^a_{eq})}{1+\delta_1})}=\frac{\partial}{\partial\sigma^a_{eq}}\frac{\int_{\frac{\delta_1 w(\sigma^a_{eq})}{1+\delta_1}}^{w(\sigma^a_{eq})}\underline{V}dF_{\underline{V}}(\underline{V})}{\lambda_2(\sigma^a_{eq})+\lambda_3(\sigma^a_{eq})}$, which by Leibniz integration rule is equal to

$$\frac{w'(\sigma^a_{eq})w(\sigma^a_{eq})\{f_{\underline{V}}(w(\sigma^a_{eq}))-(\frac{\delta_1}{1+\delta_1})^2 f_{\underline{V}}(\frac{\delta_1}{1+\delta_1}w(\sigma^a_{eq}))\}}{\lambda_2(\sigma^a_{eq})+\lambda_3(\sigma^a_{eq})}-\frac{\lambda_2'(\sigma^a_{eq})E[\underline{V}|\underline{V}\in(\frac{\delta_1 w(\sigma^a_{eq})}{1+\delta_1},w(\sigma^a_{eq}))]}{\lambda_2(\sigma^a_{eq})+\lambda_3(\sigma^a_{eq})}.$$

$\frac{\partial E[\underline{V}|\underline{V}\in(\frac{\delta_1 w(\sigma^b_{eq})}{1+\delta_1},\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2})]}{\partial\sigma^b_{eq}}=\frac{\partial}{\partial\sigma^b_{eq}}\frac{\int_{\frac{\delta_1 w(\sigma^b_{eq})}{1+\delta_1}}^{\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2}}\underline{V}dF_{\underline{V}}(\underline{V})}{F_{\underline{V}}(\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2})-F_{\underline{V}}(\frac{\delta_1 w(\sigma^b_{eq})}{1+\delta_1})}=\frac{\partial}{\partial\sigma^b_{eq}}\frac{\int_{\frac{\delta_1 w(\sigma^b_{eq})}{1+\delta_1}}^{\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2}}\underline{V}dF_{\underline{V}}(\underline{V})}{\lambda_2(\sigma)}$, which by Leibniz integration rule is equal to

$$\frac{w'(\sigma^b_{eq})w(\sigma^b_{eq})\{(\frac{\delta_1-\delta_1^2}{1-\delta_2+\delta_1-\delta_1^2})^2 f_{\underline{V}}(\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2})-(\frac{\delta_1}{1+\delta_1})^2 f_{\underline{V}}(\frac{\delta_1 w(\sigma^b_{eq})}{1+\delta_1})\}}{\lambda_2(\sigma^b_{eq})}-\frac{\lambda_2'(\sigma^b_{eq})E[\underline{V}|\underline{V}\in(\frac{\delta_1 w(\sigma^b_{eq})}{1+\delta_1},\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2})]}{\lambda_2(\sigma^b_{eq})}.$$

$\frac{\partial E[\underline{V}|\underline{V}\in(\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2},w(\sigma^b_{eq}))]}{\partial\sigma^b_{eq}}=\frac{\partial}{\partial\sigma^b_{eq}}\frac{\int_{\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2}}^{w(\sigma^b_{eq})}\underline{V}dF_{\underline{V}}(\underline{V})}{F_{\underline{V}}(w(\sigma))-F_{\underline{V}}(\frac{\delta}{1+\delta}w(\sigma))}=\frac{\partial}{\partial\sigma^b_{eq}}\frac{\int_{\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2}}^{w(\sigma^b_{eq})}\underline{V}dF_{\underline{V}}(\underline{V})}{\lambda_3(\sigma^b_{eq})}$, which by Leibniz integration rule is equal to

$$\frac{w'(\sigma^b_{eq})w(\sigma^b_{eq})\{f_{\underline{V}}(w(\sigma))-(\frac{\delta_1-\delta_1^2}{1-\delta_2+\delta_1-\delta_1^2})^2 f_{\underline{V}}(\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})\}}{\lambda_3(\sigma^b_{eq})}-\frac{\lambda_3'(\sigma)E[\underline{V}|\underline{V}\in(\frac{(\delta_1-\delta_1^2)w(\sigma^b_{eq})}{1-\delta_2+\delta_1-\delta_1^2},w(\sigma^b_{eq}))]}{\lambda_3(\sigma)}.$$

Next, we take derivatives of $\lambda_1(\sigma)=F_{\underline{V}}(\underline{V}<\frac{\delta_1 w(\sigma)}{1+\delta_1})$, $\lambda_2(\sigma)=F_{\underline{V}}(\frac{\delta_1 w(\sigma)}{1+\delta_1}<\underline{V}<\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})= F_{\underline{V}}(\underline{V}<\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})-F_{\underline{V}}(\underline{V}<\frac{\delta_1 w(\sigma)}{1+\delta_1})$, and $\lambda_3(\sigma)=F_{\underline{V}}(\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2}<\underline{V}<w(\sigma))= F_{\underline{V}}(\underline{V}<w(\sigma))-F_{\underline{V}}(\underline{V}<\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})$ and plug them into the resulting expansions.

After plugging in the resulting expansions into equations (6a) and (6b) and noting that $\lambda_1'(\sigma)=f(\frac{\delta_1 w(\sigma)}{1+\delta_1})\frac{\delta_1 w'(\sigma)}{1+\delta_1}$, $\lambda_2'(\sigma)=f(\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})\frac{(\delta_1-\delta_1^2)w'(\sigma)}{1-\delta_2+\delta_1-\delta_1^2}-f(\frac{\delta_1 w(\sigma)}{1+\delta_1})\frac{\delta_1 w'(\sigma)}{1+\delta_1}$, and $\lambda_3'(\sigma)=f(w(\sigma))w'(\sigma)-f(\frac{(\delta_1-\delta_1^2)w(\sigma)}{1-\delta_2+\delta_1-\delta_1^2})\frac{(\delta_1-\delta_1^2)w'(\sigma)}{1-\delta_2+\delta_1-\delta_1^2}$ we obtain equations (1.7a) and (1.7b).

A.2 Appendix for Chapter 2: Tests of Directions of the Bias

A.2.1 Repeated versus Non-repeated Bidders

The most important assumption of Proposition 2.1 is that bidders follow a single-stage ascending second-price auction bidding strategy in the presence of a secret reserve price. Since a bidder may expect that an item listed in an auction with a secret reserve price is very likely to be relisted again, the highest unsuccessful bid of this bidder does not necessarily reveal her valuation. We call bidders who take into account the possibility of participation in a future listing due to a secret reserve price as repeated bidders. The change in the bidding behavior of such bidders is driven by a number of factors.

First of all, a repeated bidder realizes that by participating in an ascending second-price auction she reveals information about her valuation to other repeated bidders and that other repeated bidders may exploit this information in future stages. As a result, a repeated bidder in the current stage auction may not bid truthfully. We do not intend to go into more details about this aspect of repeated bidders; for more details, see Bergemann and Said (2011).

Second, a repeated bidder realizes that her bid reveals information about her valuation to the seller, and the seller can exploit this information in the future as well. Under the assumption that the seller has imperfect information about the distribution function of bidders' valuations, a repeated bidder acting strategically has incentives to under-bid to convince the seller of a different distribution function. By under-bidding and forcing the seller to have a wrong update about the distribution function of valuations, a repeated bidder may obtain substantial gains if the seller with a wrong updated belief chooses a selling format with a higher trade surplus allocated to bidders or changes auction parameters benefitting participating bidders.

Of course, if bidders in auctions with a secret reserve price do not intend to participate in future sales, then they have no incentive to act strategically and their bidding behavior in auctions with a secret reserve price coincides with their bidding behavior in single-stage auctions with a secret reserve price. To make sure that there is no strategic component in bidders' behavior in our sample and the identification assumption is satisfied, we remove all auctions were there are two or more identical bidders participating in two or more stages of a multi-stage auction. We present the distribution of auctions by the number of identical bidders across relistings in the next table.

The results in Table A1 show that there is a sizeable share of auctions in our data, where there are at least two identical bidders participating in at least two listings of the same tractor. To satisfy the identification assumption that there are no repeated bidders in our data, we omit observations from the bottom half of Table A.1 when estimating distribution functions of bidders' valuations and when testing for determinants of secret and public reserve prices. In other words in our sample we retain only those auctions, where either all bidders in all listings are different or there is at most one identical bidder participating in at most two listings.

Type of Auctions	2-stage auctions	3-stage auctions	4-stage auctions	5-stage auctions
Auctions with at most 1 identical bidder in at most 2 listings (number)	709	76	20	5
Auctions with at most 1 identical bidder in at most 2 listings (percent)	69.04%	46.34%	38.46%	23.81%
Auctions with multiple identical bidders in at most 2 listings (number)	318	88	32	16
Auctions with multiple identical bidders in at most 2 listings (percent)	30.96%	53.66%	61.54%	76.19%

Table A.1. Share of identical bidders in different stages in repeated auctions

A.2.2 Entry Rates

To test the assumption that entry decisions of bidders do not depend on the presence of a secret reserve price and/or BINs, we regress the number of bidders on a number of determinants including the presence of a secret reserve price, the size of a public reserve price, and the size of a BIN.

The results in Table A.2 show that the presence of a secret reserve price has a strong negative impact on entry. For instance, the presence if a secret reserve price lowers the average number of bidders by more than 0.8 bidders. In addition, note that although the simultaneous presence of a BIN and a secret reserve price has a statistically significant negative impact on entry, the size of the impact is quite small. Hence, we find evidence to support our claim that the simultaneous use of a BIN and a secret reserve price does not excessively reduce entry.

Variable	Coefficient	Standard Error
Public reserve$_i$	-0.0001***	0.000001
Secret reserve dummy$_i$	-0.206***	0.004
BIN$_i$	0.000002***	0.0000007
BIN$_i$*Secret reserve dummy$_i$	-0.00001***	0.0000008
Public reserve$_i$*Secret reserve dummy$_i$	0.00008***	0.000001
John Deere$_i$	0.080***	0.003
Age$_i$	-0.008***	0.0001
Engine HP$_i$	0.002***	0.00005
ln(Seller feedback$_i$)	0.022***	0.0008
ln(Buyer feedback$_i$)	0.008***	0.0008
Constant	2.589***	0.006
Number of observations		43138

Dependent variable is the number of bidders; ***-statistical significance at 1%;

Table A.2. Poisson regression of the number of bidders

The negative impact of a secret reserve price on entry suggests that both the non-parametrically estimated distribution function of valuations and the empirical distribution of bids, which uninformed sellers use to form their beliefs about the distribution functions of valuations, are likely to over-represent higher valuations and under-represent lower valuations. It is natural to expect that a bidder with a lower valuation is less likely to enter an auction with a secret reserve price, since the probability of winning in such an auction given the equilibrium strategy of an uninformed seller is virtually zero. To test whether the presence of a secret reserve price biases the distribution of valuations toward the right tail, we regress average bids on a secret reserve price dummy. The results in Table A.3 show that the presence of a secret reserve price raises the average bid. Hence, the distribution of bids observed by an uninformed seller and our estimated distribution functions are skewed to the right.

Variable	Coefficient	Robust Standard Error
Number of bidders$_i$	26.349***	5.109
John Deere$_i$	1214.455***	91.461
Age$_i$	-81.396***	1.963
Engine HP$_i$	27.612***	1.186
Public reserve$_i$	0.785***	0.023
Secret reserve dummy$_i$	400.334***	64.826
BIN$_i$	.095***	0.012
Constant	2783.676***	95.752
Number of observations		13057
R^2		0.592

The dependent variable is the average bid; ***-statistical significance at 1%;

Table A.3. OLS regression of the average bid on a secret reserve dummy

A.2.3 Independent versus Inter-dependent Valuations

The identification of the highest valuation from the highest observable bid in unsuccessful auctions is based on the assumption that bidders' valuations are independently distributed. If bidders' valuations have a commonly distributed component, then a bidder with the highest bid over-bids relative to her valuation[1].

We argue that bidders' valuations are independently distributed, although a commonly distributed component is likely to be present as well. The main justification for the independence assumption comes from the nature of the items on sale in our sample. In section 2 we show that the average age of the tractors in our sample is about 22 years with a large share of tractors exceeding 30 years. Since most of the tractors in the sample are old, they are probably purchased for own use rather than for a later resale. By "own use" we mean that the bidders most likely intend to use the purchased tractors for farming purposes or for spare parts. In either case, the bidders' valuations of the tractors on sale are mostly guided by bidders' individual preferences and less by valuations of other bidders. Based on this consideration we argue that the independence assumption is likely to fit the data more accurately.

We follow the approach of Bajari and Hortacsu (2003) to empirically test whether bidders' valuations are independently distributed or have a commonly distributed component. The empirical test of Bajari and Hortacsu is based on the argument that if there is a commonly-distributed component in bidders' valuations then bidders should rationally lower their bids in larger auctions. Hence, Bajari and Hortacsu argue that the number of bidders should be negatively correlated with realized bids if there is a commonly-distributed component in bidders' valuations. In contrast, if bidders' valuations are completely independently distributed, then the number of participating bidders should not

[1] In the literature this phenomenon is known as the "winner's curse," and it occurs because a bidder with the highest bid takes the fact that she is the winner as a negative signal of the quality of the item on sale. Hence, in the presence of a commonly distributed component in bidders' valuations the highest observable bid exceeds the true highest valuation. For an overview, see Kagel and Levin (2002).

affect realized bids.

To control for affiliation among bidders' valuations we include dummies for 18 categories of tractors. In addition, since the number of bidders is endogenous to realized bids, we use the minimum bid as an instrument for the number of bidders. Since we regress individual realized bids, we also include bidders' feedback scores among determinants. We present results of our test in Table A.4.

The IV regression results in Table A.4 show that the number of bidders is statistically significant and negatively correlated with realized bids. According to Bajari and Hortacsu (2003), this indicates that bidder's valuations have a commonly distributed component, and the highest bids in unsuccessful auctions exceed their underlying valuations.

Variable	Coefficient	Robust Standard Error
Number of bidders$_i$	-1222.166***	198.562
Secret reserve dummy$_i$	-1689.940***	386.458
BIN$_i$*Secret reserve dummy$_i$	-3796.384***	619.197
BIN$_i$	0.073**	0.033
ln(Seller feedback$_i$)	404.160***	142.116
ln(Buyer feedback$_i$)	-146.798**	68.294
Age$_i$	-209.094***	23.806
Engine HP$_i$	56.059***	13.494
John Deere$_i$	-942.256	942.939
Number of observations		43138

The dependent variable is realized bids; ***,**-statistical significance at 1%,5%;

In the table we do not report coefficients for category dummies;

Category dummies are jointly statistically significant at 1% with p-value<0.0001

Table A.4. IV regression of realized bids on the number of bidders

A.3 Appendix for Chapter 3

To simplify the empirical testing of the sorting condition in part (c) of Proposition 2 in the rest of the chapter, we state that the expected price $p^E(\theta = \alpha)$ in Proposition 2 is an upper bound of any observable price in an electronic platform. To see this, note that $p^E(\theta(\beta \in (0,1], \gamma = 0) < \alpha) < p^E(\theta(\beta = 0, \gamma = 0) = \alpha)$ by Lemma 1.

To get a better idea of this expected price in an electronic platform $p^E(\theta = \alpha)$ and it's relation to prices in a physical platform, we derive conditions when $p^E(\theta = \alpha)$ is higher than a linear combination of prices in a physical platform with weights α and $1 - \alpha$.

Proposition

If $\frac{N-2}{N-1} > F(v)$, then $p^E(\theta = \alpha) > \alpha p_H^P + (1-\alpha) p_L^P$, and if $\frac{N-2}{N-1} < F(v)$, then $p^E(\theta = \alpha) < \alpha p_H^P + (1-\alpha) p_L^P$.

Proof:

To show this result, take a second derivative of the distribution function of the second highest order statistic $G(v) = NF(v)^{N-1} - (N-1)F(v)^N$ with respect to $F(v)$ to obtain $\frac{d^2G(v)}{dF(v)^2} = N(N-1)((N-2)F(v)^{N-3} - (N-1)F(v)^{N-2})$. Note that $G(v)$ is convex if $\frac{N-2}{N-1} > F(v)$, and $G(v)$ is concave if $\frac{N-2}{N-1} > F(v)$. Further, note that if $\frac{N-2}{N-1} > F(v)$ and $G(v)$ is convex, then $\alpha G(F_H(v)) + (1-\alpha) G(F_L(v)) > G(\alpha F_H(v) + (1-\alpha) F_L(v))$ and $\int vd(\alpha G(F_H(v))) + \int vd((1-\alpha) G(F_H(v))) < \int vd(G(\alpha F_H(v) + (1-\alpha) F_L(v)))$. This implies that $p^E(\theta = \alpha) > \alpha p_H^P + (1-\alpha) p_L^P$. A similar argument can be made to demonstrate the second part of the claim in the proposition. ■