Forward-looking Bidders in Sequential Auctions

By

Robert Zeithammer

M.A. Mathematics, B.A. Mathematics and Economics
University of Pennsylvania, 1998

SUBMITTED TO THE SLOAN SCHOOL OF MANAGEMENT IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY
AT THE
MASSACHUSETTS INSTITUE OF TECHNOLOGY

June 2003

Signature of Author: ______________________________
Sloan School of Management
May 2, 2003

Certified by: ______________________________
John D.C. Little
Institute Professor
Thesis Supervisor

Accepted by: ______________________________
Birger Wernerfelt
Professor of Management Science
Chairman, Sloan Doctoral Program

Forward-looking Bidders in Sequential Auctions

by

Robert Zeithammer

Submitted to the Sloan School of Management
on May 2, 2003 in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in Management

Abstract

At Internet auction sites like eBay, similar goods are often sold in a sequence of auctions. Buyers can therefore benefit from forward-looking bidding strategies that take into account the available information about future auctions. This dissertation develops a model of such bidding, provides both experimental and econometric evidence of the model's relevance to behavior, and explores the impact of forward-looking bidding on the seller's optimal selling strategy as well as on the overall market equilibrium.

Extending prior sequential-auction theories, the proposed model assumes that bidders know their private valuations of objects auctioned in the near future. Rational bidders "bargain-hunt" in that they bid less than they would otherwise, and early bids decrease with the private values of later objects. The model's predictions are tested in two laboratory experiments, both involving a sequence of two auctions. Both experiments show that first-auction bids decrease with the private values of the future object, but the second experiment suggests that the average decrease is smaller than predicted by the theory. An econometric analysis of eBay data finds that buyers seem to look ahead, and on average adjust their bids down as a function of their private preferences for the objects sold in the near future. They also bid less when the same item they are bidding on is available within the next few auctions.

To explore the supply-side of a sequential auction marketplace, the dissertation analyzes a model of a long-lived monopolist facing overlapping generations of forward-looking buyers. When the seller learns about the current auction-market demand from past prices, bargain-hunting poses not only the obvious cost of a lower average revenue, but it also provides a benefit to the seller by making prices more informative. In equilibrium, the seller limits the extent of bargain-hunting by threatening to withhold future supply, but the threat is only credible when the profitability of the auction-market is close to the seller's outside option. Therefore, bargain-hunting can coexist with strategic selling, but is shown to be a self-regulating phenomenon that diminishes when the existence of the auction-market is threatened by an outside spot-market.

Thesis Supervisor: John D.C. Little
Title: Institute Professor

Acknowledgements

This work would not exist without generous help from my family, mentors, and friends. I would like to thank Colette Wiseman for her patient companionship and her magical ability to put problems in perspective. I would also like to thank my parents and my brother for supporting me throughout my studies in a distant foreign country so far from home, and for always understanding and encouraging me in all my pursuits.

I benefited greatly from constant interaction with the entire faculty of the Marketing department while at MIT. I would like to thank Dan Ariely for always being candid with his advice about life, and generous with his advice about research. Jin Gyo Kim for sharing his knowledge of statistics as well as his actual estimation code that plays a key role in my research on choice models. Glen Urban for always reminding me of the real world our ivory tower is built on. John Hauser for great research advice and rigorous promotion of the scientific method. Duncan Simester for helping me position and strengthen the first chapter of this dissertation. Drazen Prelec for encouraging my interests in behavioral research. Shane Frederick for his exceptional friendliness to the doctoral students and for sharing his inexhaustible mental repository of mind-provoking puzzles and facts. Nader Tavassoli for his help during the job-market, and for giving me the opportunity to learn about teaching as his assistant. Ely Dahan for sharing his research ideas with me and helping me put together my first conference presentation.

I would also like to thank Elie Ofek for his help and advice on the second essay of this dissertation. Peter Fader for helping me discover the fascinating field of Marketing, and for guiding me through my baby steps as a researcher. Bruce Hardie for taking over where Peter left off during my stay in London, and for encouraging me to go to MIT and get a PhD.

I especially want to thank John Little and Birger Wernerfelt. John Little has taken me under his wing from the beginning of the program, and generously mentored me throughout my studies. He continues to amaze me with his critical eye for exciting new problems, his unrelenting pursuit of applicable models, and his enthusiasm for research in general. Birger Wernerfelt has opened my eyes to the delicate world of capturing complicated situations with simple models. His intuition has saved me lots of time and effort by foreseeing key tradeoffs before formally deriving them, and teaching me how to focus my energy in the most promising direction. It is not a cliché to say that Birger Wernerfelt makes the Sloan Marketing PhD program excellent. He has always been there for not just me, but for all other students, always providing sound and honest advice and active support.

The PhD students in the Marketing group have been yet another source of learning, motivation and inspiration. I would like to thank Dave Godes, Rebecca Hamilton, Kristina Shampan'er, Jiwoong Shin, Jose Silva, Olivier Toubia, Ray Weaver, and Wei Wu for being great friends and colleagues. In particular, I would like to thank On Amir for sharing his wisdom and energy throughout the whole job-market and dissertation process. We have jumped through all the hoops together, and I cannot imagine doing it without his help, encouragement, and calm. I would also like to thank Dina Mayzlin for her continued friendship and support of all the doctoral students even after her graduation and departure from Cambridge.

My life at MIT would have been incomplete without my friends. I would especially like to thank Hector Briceño for all the climbing and hiking adventures he made possible, and to Martin Zálešák for being the best housemate I have ever had. Also, the MIT Outing Club is the coolest outdoor club in the world, and I would like to thank all of its numerous members for sharing with me their enthusiasm and love of the outdoors. Finally, I am grateful to all my friends back in Czech Republic who kept in touch over the many years and the great distance. Aleš Benda, Lukáš Kohl, Martin Kovář, and Milena Kurillová have always welcomed me back home as if I had left just a week ago, I can't wait to see them all again.

Table of Contents

Chapter 1: Forward-Looking Bidders in Sequential Auctions: Theory of bargain-hunting on Bay

Introduction

Internet auctions represent a relatively large and growing sector of the economy. After starting in 1995 as a collection of unrelated auctions for collectibles, eBay has evolved into a market, in which a large number of diverse goods are traded. In 2002, over 600 million items were sold for an estimated $14.8 billion, a figure likely to keep increasing with the expected growth of the user base and an ongoing expansion into Asia.[1] The model analyzed in this paper explores certain properties of demand in such an auction-driven market.

The goods sold in Internet auctions are often similar and mutually substitutable, especially in categories involving mass-produced consumer goods like computers, electronics and CDs. The fact that mass-produced consumer goods like digital cameras are sold by auction is an interesting novelty in itself, because before the Internet, this ancient selling format had mostly been used to sell unique goods like artwork, collectibles, and mineral rights. When similar products are offered for sale in different auctions, the individual auctions compete for bidder demand, and the bidders have to adopt a multi-auction bidding strategy. For example, a consumer looking for a small digital camera will find hundreds of acceptable options on eBay, but only wants to buy one camera at the end of the day. His bidding strategy thus needs to consider both which auctions to enter, and how to bid in those auctions. The individual auctions are organized implicitly into a sequence by their ending times, so the bidder's problem reduces to considering the "current" auction (the auction that will end first) while optimally incorporating the option value of "future" auctions. This paper develops a theoretical model of optimal bidding in a sequence of auctions for similar substitutable goods, and

[1] The net income of eBay itself was $250 million, suggesting that online auctions are a permanent phenomenon since they provide a viable source of profits to the auctioneer even in times of an economic downturn and dot-com skepticism.

provides empirical evidence of the model's relevance to actual behavior both in a laboratory and on eBay.

The theoretical model highlights the trade-offs involved in optimizing the bidding strategy across two auctions, and characterizes the bidding strategies that best solve those tradeoffs. In particular, the optimal bidding strategy takes into account the existence of future auctions, leading rational bidders to "bargain-hunt", i.e. bid less than they would if there were no other future auctions. The model focuses on the realistic situation, in which the two substitutable objects sold in sequence are similar, but not necessarily identical, and consumers vary in their preferences for the objects. In the category of digital cameras, for example, each camera offers a different set of features like size, resolution, and battery-life, and different consumers prefer different bundles of features. To reflect online reality, and in contrast with most prior sequential-auction theories that assume bidder uncertainty about the future, the proposed model assumes that bidders know their preferences for objects auctioned in the near future. On eBay, for example, a usual auction lasts a week, so the bidders know with certainty which auctions will be ending in the next seven days. Unlike a world in which bidders are uncertain about objects sold in the future, optimal bidding under this new "future-knowledge" assumption incorporates personal preferences for the known future objects, forcing bidders to assess the entire known supply stream in forming a bid.

Anectodal evidence from the eBay community chat-board suggests that this abstraction to a sequence of auctions with known future items rings true with at least some eBay bidders: *"Place bids on only one item at a time and put all the rest on "watch this item". If you are outbid on first item move to the next ending time on your watching page."*[2] The eBay webpage design reinforces the sequential conceptualization by listing auctions in a sequence, with the default ordering by ending time (see Appendix 1.1). Bidders can also place future auctions on their private watch-lists, improving their ability to consider their future options.

The theoretical model places relatively high demands on bidder's ability to be forward-looking and strategic. It is therefore necessary to check whether real human

[2] An eBay newsgroup respondent to my "What do I do with multiple auctions for similar objects?" question posted in November 2000.

bidders behave similarly to the rational bidders of the theoretical model. I conducted two laboratory experiments to test the model's predictions about the properties of bargain-hunting. Both experiments involved two sequential auctions, and each controlled a different aspect of the environment between subjects. In the first experiment, I sold real products to subjects while controlling their knowledge of the future auction. The future-knowledge was controlled at three levels: no knowledge, uncertainty, and certainty of the future object. The situation was very realistic, but it did not offer a good measurement or control of the person-specific product valuations. Since the model describes the relationship between product-valuations and optimal bids, a tight test of the model's predictions has to involve control of each subject's valuations for both products. In the second experiment, I therefore induced controlled product valuations by selling hypothetical and otherwise worthless objects, and paying the bidders according to the surpluses they earned. Instead of controlling future-knowledge, I controlled the degree of substitutability of the two objects. By no longer requiring the control of future knowledge and hence not requiring some subjects to be surprised by the announcement of the second auction, I was also able to repeat the sequence many times and thus allow the subjects to learn from experience. The results of both experiments broadly confirm most theoretical predictions while exposing some limitations of the theory. In particular the subjects in the first experiments do not seem to correctly incorporate the intensity of competition manipulated as the number of other bidders bidding on the same object, and the subjects in the second experiment exhibit persistent heterogeneity in strategies not captured by the theory.

The relevance of the proposed model to actual bidding on eBay is then demonstrated using bidding data from eBay's MP3-player category. An econometric model designed to distinguish among three plausible future-knowledge assumptions provides strong support for the assumption that bidders consider near-future auctions in a way predicted by the proposed model. Not only are the bidders forward-looking, they also bid as if incorporating their relative preference for specific future items into their bids on the current item.

Literature review

This work draws on literatures studying online auctions specifically, and multi-object auctions in general. Online auctions have only existed for about seven years, so academic research focusing on them is relatively scarce. Ariely and Simonson (2001) provide a good review of the literature as well as a general framework for studying the rich set of phenomena associated with online auctions. The issue of multi-auction online bidding has not been addressed except for work by Bajari and Hortacsu (2000), who study bidder entry, and a paper by Dholakia and Soltysinski (2001) who find a "herding bias" – consumers flocking to popular auctions despite the existence of other auctions for substitute items. Outside of the game-theoretic equilibrium paradigm, several papers have tackled the problem of optimal bidding in sequential auctions by applying optimal control methodology (for a recent example, see Arora et. al., 2002). Invariably, this line of research assumes a single strategic bidder facing non-strategic competition, limiting the relevance of the conclusions and the kind of competitive phenomena that can be investigated. While some of the qualitative conclusions derived from these models extend to the more realistic world with all players being strategic, many results on the precise properties of the optimal bidding strategy do not extend, and some interesting issues, like the impact on future competition among bidders and the impact of other bidders also bidding strategically, are simply assumed away.

The theoretical model proposed here is not confined to online auctions, and contributes to the general auction theory literature. Most of the vast auction theory literature focuses on a theory of a single isolated auction, and investigations of the issues arising from multiple auctions are comparatively rare (Klemperer 1999 provides a recent review of the auction theory literature). Before outlining the contribution of the proposed model, I briefly outline the three major epochs in the literature on multiple auctions with an emphasis on sequential auctions. Besides a section in Vickrey's classic 1961 paper, the literature on multi-object auctions in general and sequential auctions in particular started with a lone paper by Ortega-Reichert who examined the deception and information-revelation in a dependent private-value setting (Ortega-Reichert 1968). Some years later, the second wave of interest in sequential issues started with the discrepancy between the empirical findings of declining prices (Ashenfelter 1989 and many others), and the then-state-of-art theoretical predictions. Declining prices are at odds with the model of sequential auctions for identical goods by Milgrom and Weber (1982b) and Weber (1983), but the more general model of stochastically-equivalent goods by Engelbrecht-Wiggans (1994) does not rule it out. Since the empirical discoveries of the price-decline "anomaly", several sequential-auction theories focused on price-trends in finite sequences of auctions. Please see Deltas and Kosmopoulou (2002) for a summary of the different possible influences on price-trends ranging from risk-aversion to non-strategic bidding. Finally, the third wave of interest in multi-object auctions was spurred by the FCC spectrum auctions in the 1990s. Motivated by the importance and the complexity of the task, several recent papers use the mechanism-design framework to investigate the efficiency and optimality of different multi-object sale-formats including simultaneous auctions and combinatorial auctions. (Ausubel & Cramton 1998, Krishna & Rosenthal 1996, Vohra & deVries 2003).

Having sketched the layout of the economic literature on multiple auctions, I can now position this paper. My model does not propose a new design of a selling mechanism, and it does not study price-trends in finite sequences of auctions. Instead I take the sequential-auction mechanism as given, and I study the influence of bidder knowledge of the objects auctioned in the near future. In particular, I propose that online bidders know their valuations of the near future objects, and I investigate the impact of this knowledge on optimal bidding.

The proposed model extends to multiple bidders the work of Gale and Hausch (1994), who examine the special case of two bidders and two auctions.[3] The extension is non-trivial because with just two bidders, the loser of the first auction has no competition in the second auction, and hence the second-object's private value is exactly the expected future surplus of the first-period loser. Therefore, Gale and Hausch elegantly but unrealistically avoid the equilibrium analysis of the relationship between current bidding strategy and future competition.

[3] Their model is also different in its assumption about zero bids in that they allow trade (by coin-toss) when all bidders bid exactly zero.

The existing models of sequential auctions with more than two bidders consider either identical objects (Milgrom and Weber 1982b), or similar but unknown future objects (Engelbrecht-Wiggans 1994). Technically, the identical-goods model of Milgrom and Weber (1982b) involves certainty of the future objects, but it confounds this knowledge with the additional assumption of identical values. Moreover, the identical-goods model is too restrictive in that the equilibrium bidding strategy depends very heavily on common knowledge among bidders that the goods are exactly identical. The proposed model is therefore the first model that separates the influence of bidder future-knowledge from the relationship between the private values of the objects sold. Because the assumption that bidders know their valuations of the near future items is natural in a typical online-auction setting, the model contributes to our understanding of demand in online auction-driven markets.

Theory

The general trade-offs involved in sequential-auction bidding can be outlined in relative generality. Suppose there are several auctions selling similar goods in a sequence, the first auction is currently open for bidding, and the rest of the auctions will occur in the future. Since winning the current item reduces the desirability of the future items to the winner, bidders face a trade-off between winning now and winning later. This trade-off makes the bidders decrease their current bids, offering to only buy the current item for a price low enough to compensate for the opportunity cost of winning. When bidders only want one item, the opportunity cost of winning the current item is exactly the expected future surplus.

When bidders know which particular items will be sold in the future auctions, they can use that information in estimating the opportunity cost of winning today, because the expected future surplus is a function of preferences for future items. When preferences are expressed as monetary valuation, the expected future surplus increases in the private valuation of future items, and so the optimal current bid has to decrease in the private valuation of future items.

When all bidders have and use such information about the future, they also need to account for the implied change in future competition. Future competition may increase because the bidders who are willing to pay more for the future items are now more likely to lose the current auction compared to the bidders who would pay more for the current item. However, if the bidders' private valuations of the current and future objects are correlated across time, future competition may decrease because bidders with highest overall values are more likely to win the first auction. The resulting pure-strategy equilibrium is characterized by an expected surplus function that takes all the above considerations into account.

Having outlined the main trade-offs considered by the theoretical model, I can discuss the key assumptions. To obtain a tractable model of strategic bidding in multiple auctions, it is necessary to assume that the component individual auctions are relatively simple. The key simplifying assumption about the individual auctions is that bidders have independent valuations for the objects. This assumption is non-trivial in that it excludes many of popular eBay product-categories, namely antiques and collectibles, in which valuations are most likely "affiliated" across bidders (Milgrom & Weber 1982a). On the other hand, mass-produced goods like consumer electronics are usually purchased for private use with resale opportunities limited by rapid obsolescence, thus justifying the assumption.

The second assumption simplifies the online auction dynamics and anchors the auctions in a one-dimensional world measured by time. Online auctions usually remain open for several days, thus creating room for strategically rich within-auction behavioral dynamics. I abstract from these within-auction dynamics, and model each individual auction as an instantaneous sealed-bid auction occurring at the time of the actual auction's end. This reduction of each online auction into a sealed-bid auction is supported by the fact that bidding on eBay tends to happen at the very end of each auction without giving the competitive bidders time to react (Roth & Ockenfels 2002). Since ending times of online auctions are not synchronized, the abstraction results in a model of sequential auctions. Note that abstracting away from the within-auction dynamics is relatively more reasonable in the assumed independent-value world, where bidder's do not try to learn about their own preferences from other people's bids.

Given the above assumptions, I first solve the model with two auctions and a simple distribution of preferences, and then discuss possible generalizations.

Model

Two sequential second-price sealed-bid auctions 1 and 2 each sell one item to a group of N risk-neutral unit-demand bidders. Each bidder knows his private single-item valuations (v_1,v_2) at the start of the game, and the valuations are drawn independently across bidders from the uniform distribution $f(v_1,v_2)=1$ on $[0,1]x[0,1]$ [4]. The objects are similar in that $f(u,v) = f(v,u)$, with a marginal cumulative distribution $F(v)=v$. The bidders are only interested in winning one item, so the winner of the first object does not bid again.

[4] Note the compact support of f.

Assume the two auctions are conducted promptly right after each other, so there is no discounting of the second auction's outcome.

Optimal first-period bidding strategy

In the second auction, it is a dominant strategy to bid your value v_2. Suppose somebody bid more than zero in the first period. Then, since the first-period winner does not bid again in the second auction, there are N-1 bidders in the second auction. The N-1 first-period losers bidding their valuations give rise to the future-surplus function that captures the expected surplus available to losers of the first auction in the second period as a function of the second-period valuation. The surplus function satisfies the following equation 1.1:

$$S(v_2) = \int_0^{v_2} (v_2 - c_2) dG_2(c_2 \mid S) = \int_0^{v_2} G_2(c_2 \mid S) dc_2 \qquad (1.1)$$

Where c_2 is the second-period highest competitive bid, i.e. the maximum c_2 among N-2 other bidders who lost the first auction, and G_2 is the distribution of c_2. The second equality in (1) follows from integration by parts. It will become evident that first-period bidding, and hence the determination of first-period losers, depends on S as suggested by the conditioning of the competitive-price distributions on S. Therefore, equation (1.1) defines the future-surplus function only implicitly.

Given the second-period surplus, the bidders solve the following problem in the first period:

$$b_1(v_1, v_2) = \underset{\beta>0}{\operatorname{argmax}} \left\{ \int_0^{\beta} (v_1 - c_1) dG_1(c_1) + \int_{\beta}^{1} S(v_2 \mid c_1) dG_1(c_1) \right\} \qquad (1.2)$$

Since the bids in the first auction are in general a function of v_2's and c_1 is an upper bound on others' first-period bids, $S(v_2)$ needs to be conditioned on c_1 as shown in equation (1.2). By definition of c_1 as the highest competitive bid, when $\beta < c_1$ then c_1 is the winning bid, and further conditioning the surplus by β is not needed (a key simplification made possible by independence of bidders of each other). Therefore, the following first-order condition (FOC) characterizes optimal bidding in the first period:

$$0=[v_1 - \beta - S(v_2 \mid \beta)]g_1(\beta) \Rightarrow b_1 = \max[0, v_1 - S(v_2 \mid c_1 = b_1)] \tag{1.3}$$

where the g_1 is the pdf of c_1, and the bid is constrained to be non-negative by definition. The first-order condition provides interesting intuition. When evaluating the option value of the second auction, the bidder assumes that he loses the first period to a competitive bid that exactly matches his first bid. This is because that is the only situation, in which raising the first bid slightly changes the outcome of the game, and $S(v_2,b_1)$ is therefore the only opportunity cost relevant at the margin.

The corresponding second-order condition (SOC) which ensures that the FOC describes a local maximum of the bidder's expected utility boils down to $\frac{\partial S}{\partial c}(b_1) > -1$.[5] SOC is also the necessary and sufficient condition for the first bid to decrease in v_2:

$$\text{FOC} \Rightarrow \frac{\partial b_1}{\partial v_2} = \frac{-\frac{\partial S}{\partial v_2}}{1+\frac{\partial S}{\partial c}} \text{ and because } \frac{\partial S}{\partial v_2} > 0 \text{ always, } \frac{\partial b_1}{\partial v_2} < 0 \Leftrightarrow \frac{\partial S}{\partial c} > -1 \Leftrightarrow \text{SOC} \tag{1.4}$$

When each bidder's values are independent across periods and the distribution is relatively flat like the assumed uniform distribution on [0,1]x[0,1], the SOC actually implies (through the slope of bids in v_2) that S has to increase in c_1: the conditioning of S has the form of an upper bound on others' first bids, and since first bids decrease in v_2 and the values are independent across periods, relaxing the constraint makes lower competitive v_2's more likely. Therefore, $\frac{\partial S}{\partial c} > 0$ everywhere.

Whether the SOC holds fundamentally depends on how the underlying distribution of valuations plays out in the equilibrium expected-surplus function. In particular, strange things may happen when f is negatively correlated or locally very

[5] SOC: $-\left[1+\frac{\partial S}{\partial c}(\beta)\right]g_1(\beta) + [v_1 - \beta - S(v_2 \mid \beta)]\frac{\partial g_1}{\partial \beta} < 0$

steep. I assume SOC holds under the uniform assumption, and later verify that the equilibrium S indeed satisfies $\frac{\partial S}{\partial c} > 0$ and hence the SOC.

Figure 1.1: Unique solution to the first-order condition

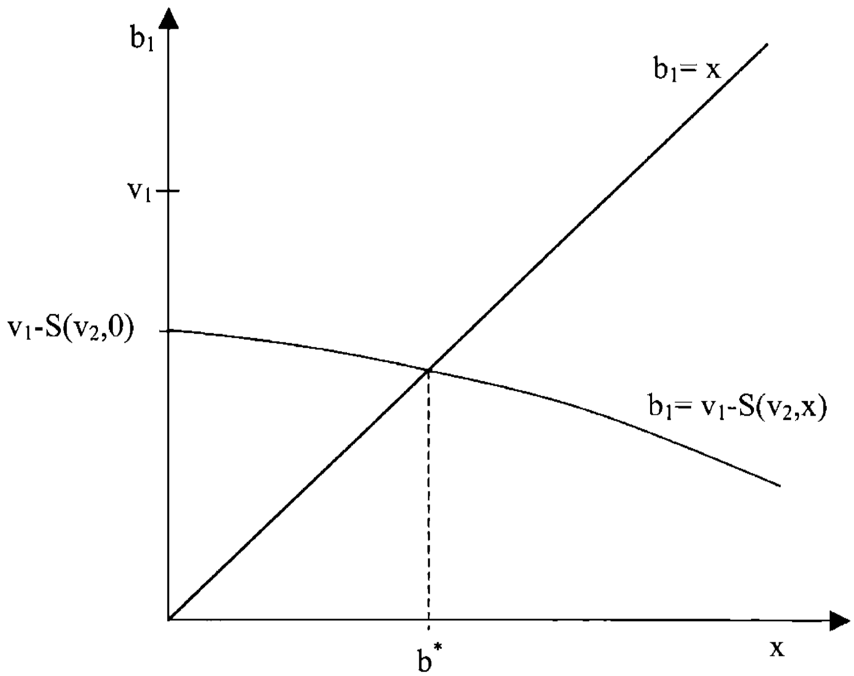

When $\frac{\partial S}{\partial c} > 0$ everywhere, the FOC has a unique solution because S is continuous and non-increasing in c_1, so if v_1-S(v_2,0) $\geq$0, the Intermediate Value Theorem implies that there is always b^* that satisfies the FOC, and the fact that S is monotonic in c_1 implies that there is a unique b^* (see illustration in Figure 1.1). Whenever v_1-S(v_2,0) is negative, the bidder bids zero, effectively waiting for the second period, so then b^*=0. Let the first-period bid therefore be a function $b_1(v_1,v_2) = b^*$.

Equilibrium existence

Suppose the bidders use some function S_B to form their first-period bids. Then, unfolding equation (1.1) produces the expected surplus function S_E given the bidding strategy:

$$S_E(v_2, c_1) = \int_0^{v_2} G_2(c_2 \mid c_1) dc_2 = \int_0^{v_2} \Pr\left(\max_{i=\{N-2 \text{ other losers}\}} \{ v_{2,i} \mid b_{1,i}(v_{1,i}, v_{2,i}) \le c_1 \} \le c_2 \right) dc_2 \qquad (1.5)$$

where $b_{1,i} = \max[0, v_{1,i} - S_B(v_{2,i}, b_{1,i})]$ for all i.

Equilibrium existence hinges on the existence of S that satisfies: $S_E=S=S_B$. In particular, the expected-surplus calculation has to be correct in equilibrium given that all bidders use the same expected-surplus function in forming their first-period bids. When the valuations are distributed uniformly on [0,1]x[0,1], such S exists, and it is increasing in c_1:

Theorem 1: If f is uniform distribution on [0,1]x[0,1], then there exists a continuous function S increasing in c such that if first-period bidders bid $b_1=\max[0,v_1- S(v_2,b_1)]$ and the surviving second-period bidders bid v_2 , then $S(v_2,c)$ is the expected second-period surplus of a first-auction loser given that the highest bid in the first auction is c. (proof in Appendix 1.2)

While analytically untractable, the equilibrium expected surplus function has several properties outlined in the following propositions:

Proposition 1.1: Given regularity conditions sufficient to ensure double differentiability of G_2, the future-surplus function S is positive, nondecreasing in both arguments, and convex in v_2. (proof in Appendix 1.2)

The expected surplus increases in the valuation because a higher valuation makes winning more likely and also increases the actual surplus conditional on winning. Since these two at-least-linearly-increasing components effectively multiply to produce the expected surplus, the convexity results. The function also increases in the highest competitive bid c because as c increases, relatively smaller competitive v_2's are more likely, and the competition hence weakens. It is important to note that the results do not depend crucially on the assumption about the number of bidders and the fact that one of the bidders drops out after the first period. Bidders still bargain hunt as long as they have a chance of winning the second item regardless of the particular assumption about the evolution of the bidder pool, be it drop-out or replenishment other bidders. As long as some bidders carry on to the second period, a necessary condition for bargain-hunting, the equilibrium considerations about the expected-surplus function remain qualitatively the same as discussed above. The overall number of bidders does effect bargain-hunting qualitatively:

Proposition 1.2: S(x,c) converges to zero pointwise exponentially in the number of bidders N for all x s.t. $F(x)<1$. (proof in Appendix 1.2)

This is an important result delineating a boundary of the phenomenon. As the number of bidders increases, it is harder and harder to win the second auction while the winner's surplus also shrinks. Therefore, the expected surplus from the second auction shrinks to zero and the two auctions become effectively isolated.

Proposition 2 does not necessarily show that bargain-hunting is impossible on eBay because there are many registered eBay users. The key parameter is not how many potential bidders there are, but how many bidders actually participate in an average auction. Bargain-hunting disappears in the unrealistic situation when all bidders participate in all auctions, but when bidders only choose small subsets of auctions to participate in, perhaps because they incur a participation cost, the economically-relevant quantity represented by "N" in the present model can be better interpreted as the ratio between demand and supply, not as the overall size of demand.

Finally, it is easy to see that S(v,c) is bounded from above by M(v), where M(x) is a "myopic" surplus function resulting from facing N-2 bidders bidding according to the distribution F, i.e. bidding their valuation. Moreover, S(v,c) approaches M(v) as c approaches 1. This is because as c increases, the bound on the competitors' first bids weakens all the way to a point when it becomes not binding. Note that in the uniform case, the expected surplus from a single isolated auction is easy to compute because the cdf G(c) of the competitive bid from N-2 bidders is just c^{N-2}. Therefore,

$$M(v) = \int_0^v G(c)dc = \int_0^v c^{N-2}dc = \frac{v^{N-1}}{N-1}.$$

Numerical approximation of the future surplus function

Given any particular f, it is possible to numerically approximate the S function by successive iterations of the relationship between second-period surplus and first-period bidding shown in equation (1.5). Starting with a guess of some S_0, first-period bids can be simulated and the resulting conditional $G_2(c_2|c_1)$ distribution approximated, leading to an improved guess of S. In other words, successive iterations of the mapping T of the Proof of Theorem 1 (see Appendix 1.2) should lead to a fixed-point S that satisfies equation (1.1). The process seems to converge fast under a wide variety of distributional assumptions.

Assuming three bidders with iid uniform (0,1) distributions, the approximated equilibrium bidding function is shown in Figure 1.2 below.

Figure 1.2: First-period bidding function, N=3, F=Uniform[0,1], (v_1,v_2) independent

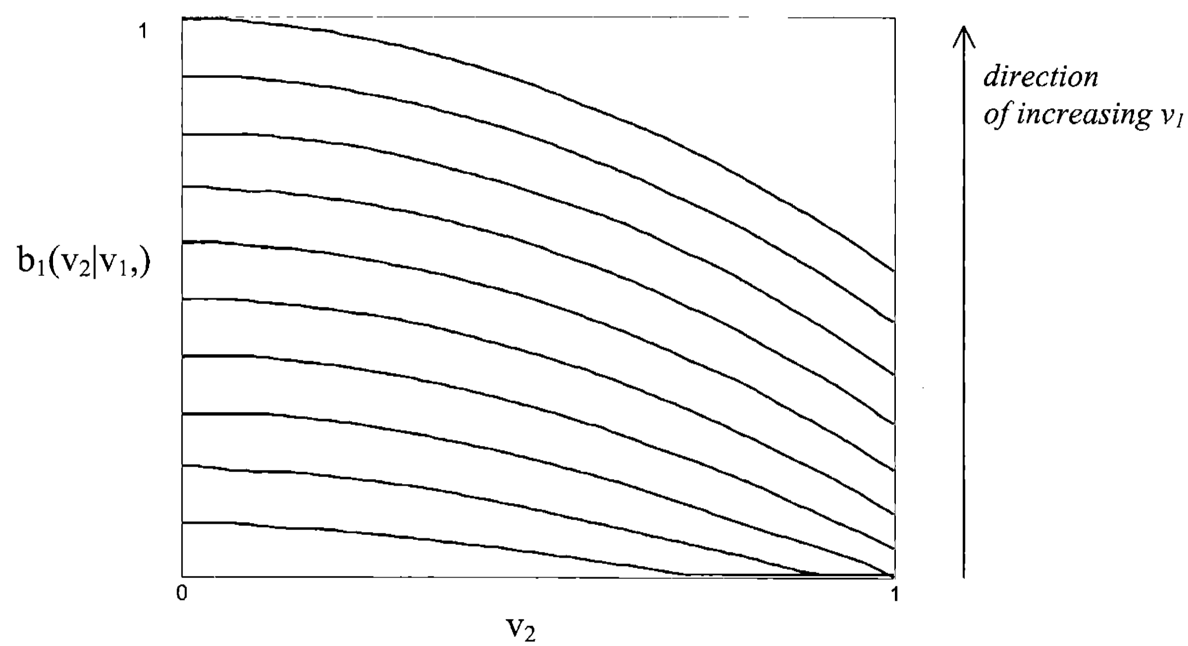

Each line in Figure 1.2 represents the change in the first-period bid with v_1 fixed and v_2 changing while the y-intercept of each curve is exactly v_1 because $b_1(0|v_1) = v_1$. As predicted, S(v,c) is non-decreasing in c, so the lines become less steep for smaller v_1's, so the less a bidder likes the first object, the less he will shade his bid for a given second object with a value of v_2. Given the fact that $b_1(v_2|1)$ is close to 1-$M(v_2)$, this effect of v_1 on bid-shading means that bids of zero are actually a lot less likely than might be expected from a naïve calculation of the second-period surplus using M(x). On the other hand, Figure 1.2 clearly shows that some waiting for the second period (bidding zero in the first) still occurs.

Another way to illustrate the equilibrium is to consider the equilibrium surplus function. Please see Figure 1.3 below for the expected surplus function at eleven levels of c_1, namely {0, 0.1,0.2...1}.

Figure 1.3: Equilibrium expected-surplus function

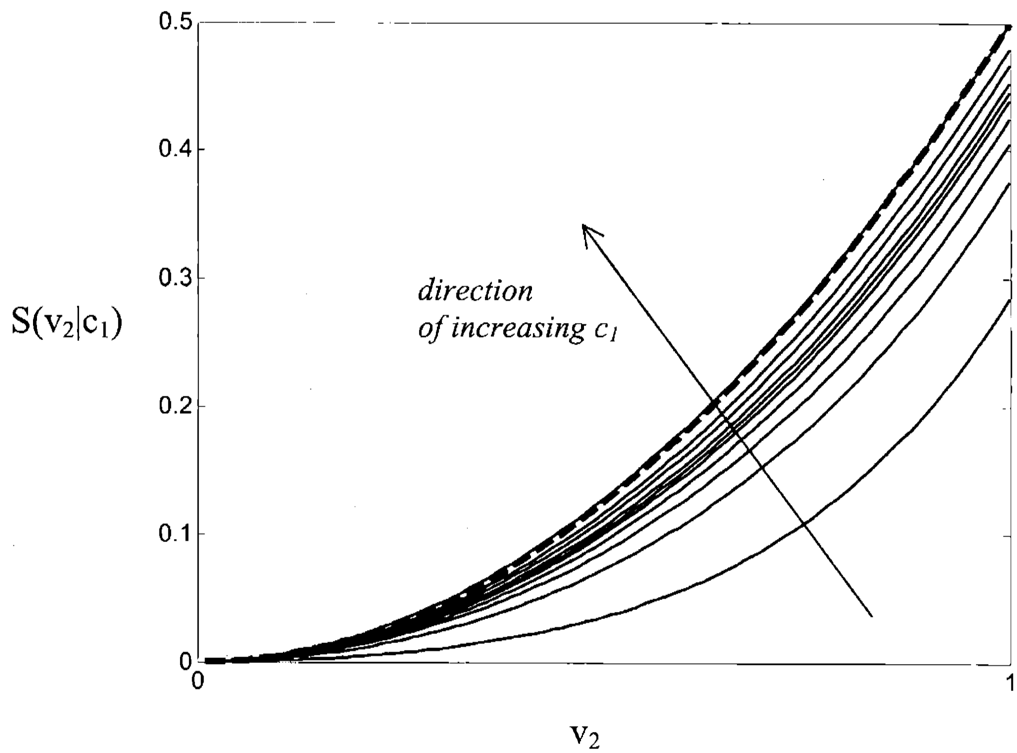

The dashed wider line in Figure 1.3 is the $M(v_2)$ surplus function. As predicted, it is identical to $S(v_2|1)$.

Relaxation of the distributional assumption

It is clear from the above discussion that a well-behaved equilibrium surplus function must be non-decreasing in c_1. This property both ensures that the first-order condition characterizes optimal bidding, and guarantees that the bids decrease in v_2 with a slope less than unity. From the proof of Theorem 1, it can be deduced that two aspects of the uniform distribution make the existence of such a well-behaved surplus function possible: independence of valuations across periods, and the fact that the uniform distribution is not locally very steep. A generalization of the distributional assumption to a class of smooth equicontinuous distributions with a low-enough bound on the slope should therefore be possible in the sense that the essence of the proof would still work. When, however, the correlation across periods gets too large or the distribution f locally very steep, strange things may happen, bids may start increasing in v_2, or an equilibrium bidding function may not exist at all.

The first-period bids decrease in v_2 on average, fundamentally because the valuations are independent across periods. A comparison with a constrained model reveals that the relative level of second-period competition crucially depends on the correlation across periods. When v_1 and v_2 are completely uncorrelated, the second-period competition is higher than random because the bidders with relatively higher v_2 are relatively more likely to lose the first auction. In particular, $S(x,c) < M(x)$. When, on the other hand $v_1 = v_2$, then, the identical-good model of Milgrom&Weber (1982b) applies, and the second-period competition is lower than random because the highest overall bidder wins the first auction.

Generalization to multiple periods

The two-period model is sufficient for illustrating the basic tradeoff between winning now and winning later present in any number of sequential auctions. When there are more than two auctions in a sequence, the bidders will still bargain-hunt in expectation of a positive future surplus. However, the optimal bidding strategies are no longer symmetric, and no longer depend only on each bidder's private information, because there is useful information in prices and other experiences encountered along the way.

To see how an outcome of earlier auctions influences bidding, we have to consider at least three auctions and focus on the middle auction. The tradeoff between winning now and winning later still persists, but the actual outcome of the first auction also matters because it is informative about the level of future competition. In particular, if first-period bids decrease in all future private valuations, then lower first-period prices mean stiffer competition in the future. At the conclusion of the first period, bidders update their information about their current and future competition, and use this information in forming second-period bids. To make matters even more complicated, the price-determining bidder now has slightly different information about the remaining competition than the bidders whose first-period bids were less than the price. This asymmetry escalates with each period. It clearly matters exactly what information about the outcome of the first auction is made public. The only information that would maintain informational symmetry among bidders would be the bid of the winner. However, revealing the winner and the price is much more realistic. Finally, since first-period trade is not assured, the bidders have to use an entirely different expected-surplus function when there is no first-period price.

Despite the above complications, qualitative arguments can still be made about the optimal second-period strategy in the three-period game. Let I_2 be the information available to the bidder in the beginning of the second period, namely the price of the first object p_1 and own first-period bid b_1 (where the bid is important only to determine whether $b_1 = p_1$ or $b_1 < p_1$). Because I consider second-price auctions, the information I_2 about second-period competition is of no use to the bidders in forming the second-period bid. To demonstrate this claim, let S_3 be the expected surplus in the third auction contingent on trade in the second period and note that the bidder solves the following problem in optimizing the second-period bid:

$$\max_b \int_0^b (v_2 - c_2) dG_2(c_2 \mid I_2) + \int_b^1 S_3(v_3 \mid I_2, c_2) dG_2(c_2 \mid I_2) \qquad (1.6)$$

So long as the distribution G_2 of the second-period competitive bids c_2 is continuous, the problem has a unique solution, and the optimal bid b does not depend on G_2:[6] $b_2 = \max[0, v_2 - S_3(v_3|I_2, b_2)]$. However, information about the third-period competition is useful because it determines the opportunity cost of winning the second object. Therefore, I_2 only enters S_3, and the first qualitative question is how. I focus on the price component of the information and assume the valuations are independent across periods.

Analogously to the situation in equation (1.1), integration by parts reduces the computation of S_3 to the evaluation of third-period equilibrium competitive bids c_3. Since third-period equilibrium competitive bids are just the maxima of N-3 third-period valuations[7], it is enough to study $F(v_3, I_2)$. Since the future potential surplus S_2 available to the first-period bidders most likely increases in v_3, first-period bids decrease in v_3, and therefore the contingent third-period valuations decrease in p_1, in turn increasing the surplus available to third-period bidders. In particular, if $s>t$, then $F(v_3, s) > F(v_3, t)$ for every v_3, so $S(v_3, s) > S(v_3, t)$. To conclude the string of cause and effect, second-period bids should therefore decrease in first-period prices. It is not clear how this preliminary conclusion would generalize to a situation with dependent values, either correlated within a bidder across periods or correlated across bidders.

Comparison with another model of sequential auctions

There are other possible specifications of the game, in which two similar objects get auctioned sequentially. Consider keeping the basic setup the same while varying how much the bidders know about the second auction. Besides the trivial myopic model that assumes the bidders have no idea about the second object, the dominant model existing in the literature considers a situation in which the second object's value gets revealed only after the first auction. This model was analyzed in detail by Engelbrecht-Wiggans (1994). He finds an iterated conditionally-dominant strategy in which remaining bidders bid

[6] Note that when deriving the optimal second-period bid, the second-period competitive bid c_2 plays a role analogous to the c_1 in the two-period model.

[7] "competitive bids" are the random prices to which competition is reduced, and "third-period valuations" are valuations of people who advance all the way to the third period.

valuation in the second period, and deflate their first-period bids from valuation by bidding $b(v_1)= v_1-E$, where E is the expected surplus from the second period. Engelbrecht-Wiggans(1993) shows that this E is exactly the difference in total social surplus (i.e. the expected value to the highest bidder) resulting from one more bidder: $E=E[\max(v;N-1)-\max(v;N-2)]$. In the case of the uniform[0,1] distribution of values, E is especially simple: $E=1/(N^2-N)$. Therefore, uncertainty about the future products makes the bidders deflate their bids by a constant, by construction independent of the second object's valuation.

Since it is possible that bidders on eBay and elsewhere in the real world account for the existence of future surplus but do not account for the particular products sold in future auctions, the stochastic-equivalent model should remain a candidate for explaining forward-looking bidding behavior in sequential auctions.

Discussion of the theory

The model illustrates how optimal bidding strategy solves the tradeoffs involved in sequential-auction bidding, and how such optimal bidding plays out in the equilibrium of the game with multiple bidders. In general, bidders take into account their own private value v_2 of the future item, and decrease their first bids monotonically in v_2.

When deciding how much to decrease the first bid below their valuation, bidders take into account the fact that should they lose the first auction, their ex ante expected surplus from the second auction changes depending on how much they bid on the first item. To compute the option value of losing the first auction, it is optimal for every bidder to assume that he loses the first period to a competitive bid that exactly matches his first bid. This is because that is the only situation, in which raising the first bid slightly changes the outcome of the game, and when one bids b_1, $S(v_2|c_1=b_1)$ is therefore exactly the option-value relevant at the margin. The pivotal nature of a second-price auction thus comes through even in a sequential context, albeit in a convoluted way. The expected surplus $S(v_2)$ depends on the highest competitive first-period bid precisely because first-period bids decrease in v_2, and lower competitive bids thus place tighter constraints on the possible valuations of the competitors who will remain in the game until the second period. Therefore, S increases in c_1.

In equilibrium, bidders also have to take into account the increase in expected future competition caused by other bidders also decreasing their bids in their own second-period valuations. Therefore, the expected-surplus function has to be correct in assessing the option value of the future given that everybody uses the very same function to form their first-period bids. Existence of such a function under the model's distributional assumptions is guaranteed by the continuity and compactness of the situation.

The proposed model employed a simple uniform assumption about the distribution of valuations independent across periods and bidders. While the independence across bidders is a key assumption, the independence across periods can be somewhat relaxed towards the more realistic positive correlations. As long as the joint distribution of valuations remains continuous and nowhere too steep, all qualitative properties of the simple uniform model should hold. When, however, the correlation gets large and approaches unity, the model's predictions do not necessarily generalize, and an equilibrium bidding function may not even exist.

Changing the assumption about the bidders' knowledge of the future auctions produces at least two plausible alternative models of behavior in sequential auctions. When the bidders have no private information about the particular items sold in the future auction, they still bargain-hunt, but their bids can no longer decline in their own private value v_2 of the future item. To highlight the difference in the future-knowledge assumption, I call this first alternative model the "uncertain" model in the rest of the paper, while the model proposed in this paper is called "certain" from this point on. When the bidders do not know at all about the future or do not consider the future, they behave myopically and bid their valuations. This trivial benchmark model will be called "myopic" throughout the rest of the paper.

Experimental Evidence

Introduction to the experiments

The theoretical model places relatively high demands on bidder's ability to be forward-looking and strategic. It is therefore necessary to check whether real human bidders behave similarly to the rational bidders of the theoretical model. I conducted two laboratory experiments designed to test two different aspects of the theory. The first experiment directly tests the importance of the new future-knowledge assumption and the impact of the number of bidders on the magnitude of bargain-hunting. The strength of the first experiment is in that it actually sells a real product to the subjects, while the weakness of the first experiment is its lack of control of the subject's private valuations. Without control of the valuations, it is hard to test for a specific functional form of the bidding function. The second experiment addresses this main weakness of the first experiment by explicitly controlling each subject's valuation in the sense that hypothetical goods are sold and the subjects are rewarded according to the surplus they earn. Also, each experiment strikes a different balance between internal and external validity, and together, the provide strengthened evidence for the effect under study. I now discuss the two experiments in turn.

Experiment 1: Real product, future-knowledge controlled

Theory summary and hypothesis development

To derive testable predictions, it is enough to consider a stylized two-period world. As shown in the theory section, the first bid (i.e. the bid on the first object) depends crucially on the bidder's knowledge of the future. Table 1.1 outlines the relationship between future knowledge and optimal equilibrium bidding strategy of a bidder with valuations (v_1,v_2). Please see the theory section of the paper for details and derivations.

Table 1.1: Equilibrium Bidding Strategy of a Bidder with valuations (v_1,v_2)

Future Knowledge	***First Bid***	***Second Bid***
Myopia/ no knowledge: *second auction no effect*	v_1	*winner:* 0 *loser:* v_2
Uncertainty: *Uncertain second value*	v_1-E	
Certainty: *Known second value* v_2	b=Max[0, v_1- S(v_2,b)]	

In Table 1.1, E is by construction a positive constant and S is the expected future-surplus function. The key property of S is that it increases in v_2 and b, and the first bid thus decreases in v_2 at an increasing rate (see Figure 1.3 for an example). The three different bidding functions are illustrated in Figure 1.4. The second-period bidding function is the same in all three cases.

Several predictions can be discerned from Table 1.1 and Figure 1.4. First, bidders who know about the future auction should bid less that those who do not know. Second, the "certain" bidders, who know their valuations of the second object (v_2), should decrease their first bid more whenever their valuation of the second object is higher. On the other hand, the "uncertain" bidders who know about the existence of the second auction should not exhibit such a relationship between the bid decrease and their yet-unknown v_2. Finally, it follows from the theoretical discussion in the previous section that all of the above effects of the future should diminish as the number of bidders in an auction increases. These theoretical predictions lead to the following experimental hypotheses:

Experimental hypotheses

H1: Knowledge of a future auction decreases the first bid

H2: In the certain condition, bidders deflate the first bid more when they prefer the second object[8].

H3: The magnitude of the H1 and H2 effects diminishes with the number of bidders participating in the auction.

Figure 1.4: Relationship between the bid on the first object and future knowledge

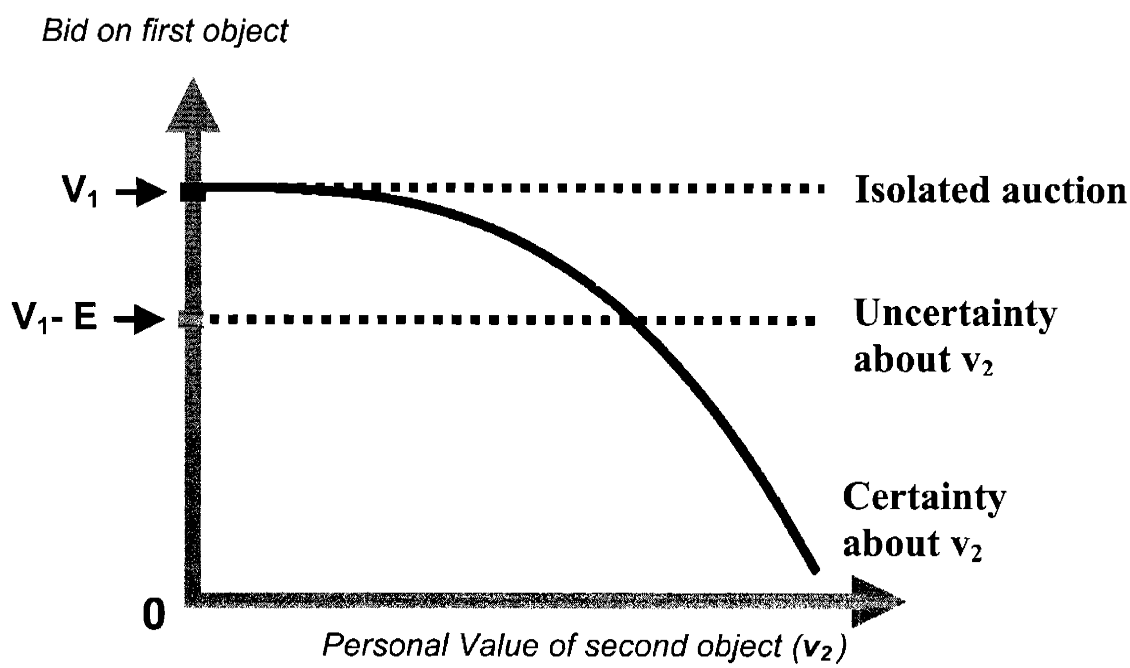

[8] For internal consistency, it is important to also check that in the uncertain condition, bidders deflate the first bid by a constant wrt personal preferences for the second object.

Experiment design

In the experiment, I sold two similar private-value substitute products (small ¼ lb. boxes of Godiva chocolates that had to be partially eaten immediately after being received) in two subsequent auctions immediately following each other. While two boxes of chocolates are not perfect substitutes as the two objects in the theory, saturation should make them substitutes and the requirement of eating some of the first box was supposed to reinforce this notion. The subjects were mostly MBA students, and they paid for the chocolates out of their own pocket.[9] Before bidding, a measure of familiarity with Godiva chocolates was collected on a five-point scale. To create a situation when the two products are similar but not identical, the two boxes were the same size, shape, color, and weight, but the first box sold was all milk chocolate, and the second box was all dark chocolate. This manipulation also enabled a measurement of relative preference for the two products: after both auctions were over, the subjects were asked to indicate whether they preferred milk or dark chocolate. About 40 percent of the subjects reported a preference for dark chocolate.

In a 2x3 between-subjects design, subjects were randomly assigned (by receiving different sets of written instructions) to six conditions arising from all combinations of two group-sizes (four or eight bidders per auction) and three future-knowledge conditions:

<u>Control</u>: the fact that there is a second auction is revealed only after the first auction.
<u>Uncertain</u>: the existence of the second auction is revealed at the outset, but the exact description of the second object remains undisclosed.
<u>Certain</u>: the existence of the second auction is revealed at the outset together with the exact description of the second object.

[9] In one class (about 50 observations), the payments were actually not collected in the end since the instructor teaching the class insisted on purchasing the chocolate himself. However, the subjects believed throughout the experiment that they would have to pay.

Results

The experiments were conducted in several different MBA classes, and there were a total of 247 observations of the first bid (13 subjects probably did not understand second-price auctions and put in very high bids; their data was excluded from the analysis)[10]. I now test the experimental hypotheses starting with the main effect.

Hypothesis 1 predicts that the knowledge of a future auction decreases the first bid. To test this hypothesis with maximum precision, I pool across the uncertain and certain conditions. The average bid in the control condition is \$4.63, while the average bid in the two conditions with future knowledge is \$4.03.[11] The difference of \$0.60 is marginally significant according to a one-tailed t-test (p=0.06). The difference in median first bids is a sizeable \$1.40. I can also test whether the two distributions (control and future-knowledge) are different using the non-parametric Wilcoxon test. The difference is marginally significant (p=0.14). Breaking apart the future-aware subjects into their respective experimental conditions does not shed further light on Hypothesis 1 since bids in both future-knowledge conditions are lower than the bids in the respective control condition. (please see Table 1.2). In summary and as predicted, the effect of future knowledge on the first bid is negative, fairly large (about 15 percent), and marginally significant.

Table 1.2: First bid by experimental treatment

	Small Group	**(N=4)**	**Large Group**	**(N=8)**
First Bid	Mean	*Error*	Mean	*Error*
Control	5.31	*0.56*	4.09	*0.40*
Uncertain	3.97	*0.43*	3.87	*0.48*
Certain	4.33	*0.40*	3.94	*0.43*

[10] The excluded bidders bid more than 15 dollars for a box of chocolates worth about \$7 in the store and receiving a median \$4 bid in this experiment. Including the data would not change the direction of any effect, it would only add noise to the data. Please see Appendix 1.4 for the analysis on all data.

[11] Note that the bids in the certain condition are actually a little bit higher on average compared to the bids in the uncertain condition. Given that more people like milk chocolate, this difference would be expected if the instructions in the uncertain condition actually make the subjects expect another box of milk chocolates.

Hypothesis 2 essentially predicts the expected surplus function to have a positive slope in the certain condition. To compute this crucial slope, I use the reported preferences for dark chocolate as a measure of relative preference for the second auction. The relevant data, along with the model, is shown in Table 1.3:

Table 1.3: Expected surplus function

	Prefer milk		***Prefer dark***	
First Bid	Mean	*Error*	Mean	*Error*
Control	$V_{milk} = 4.82$	*0.48*	$V_{dark} = 4.43$	*0.48*
Uncertain	$V_{milk} - E_m = 4.18$	*0.45*	$V_{dark} - E_d = 3.60$	*0.45*
Certain	$V_{milk} - S_m = 4.75$	*0.40*	$V_{dark} - S_d = 3.07$	*0.30*

In Table 1.3, $S_m=E[S(v_2)|\text{prefer milk}]$ and $S_d=E[S(v_2)|\text{prefer dark}]$. Hypothesis 2 essentially predicts that $S_d > S_m$, and internal consistency further predicts that $E_d = E_m$. The data in Table 1.3 imply that $S_d - S_m = \$1.28$ (s.e. 0.85). While this difference is only marginally significant ($p=0.13$), it is a very large effect, and the low significance is caused merely by the need to take a difference of differences. As expected, the implied $E_d - E_m$ is -0.19, not statistically different from zero.

Another way to measure $S_d - S_m$ is to impose the model and estimate a linear regression to get a least-squares estimate including all the data (above analysis ignores the data from the uncertain condition when estimating certain slope). I specify the linear regression such that $S_d - S_m$ and $E_d - E_m$ are coefficients to be estimated. Let $S= S_d - S_m$, $E= E_d - E_m$, $V= V_{dark} - S_m$, and $M= V_{milk} - V_{dark}$. Then, the theoretical model has the following predictions:

Table 1.4: Decomposition of effects into linear components

	Prefer milk	Prefer dark
Control	$V + M + S_m$	$V + S_m$
Uncertain	$V + M + S_m - E_m$	$V + S_m - E - E_m$
Certain	$V + M$	$V - S$

Therefore, a linear regression of the first bid on the following variables can measure both S and E directly as coefficients:

- *prefer_darkXcertain* interaction dummy (coefficient will be an estimate of S)
- *prefer_darkXuncertain* interaction dummy (coefficient will be an estimate of E)
- *prefer_milk* dummy (coefficient will be an estimate of M)
- *not_certain* dummy (coefficient will be an estimate of S_m)
- *uncertain* dummy (coefficient will be an estimate of E_m)

The estimation results are in Table 1.5 below:

Table 1.5: Regression analysis of expected surplus function (N=247, R^2=0.04)

Variable	Parameter	SE	t-value
constant	4.574	0.620	7.383
certainXprefer_dark	-1.500	0.811	-1.850
uncertainXprefer_dark	-0.398	0.808	-0.493
not certain	-0.041	0.559	-0.073
prefer_milk	0.206	0.537	0.383
uncertain	-0.537	0.588	-0.913

The regression results strengthen the above result from simple means by estimating the two crucial "slopes" to be S_d - S_m=\$1.50 (p=0.06) and E_d - E_m = 0.39 (p=0.62).

The linear regression framework can also be used to control for other variables that influence the first bid. Two important controls are the subject's bid on the second object (proxy for overall willingness to pay for chocolate), and reported familiarity with Godiva chocolate. The results are in Table 1.6 below, and they further strengthen the evidence for Hypothesis 2 by finding S_d - S_m=\$1.43 (p=0.05) and E_d - E_m = 0.18 (p=0.80).

Table 1.6: Regression analysis of expected surplus function (N=247, R^2=0.25)

Variable	Parameter	SE	t-value
constant	1.815	0.782	2.322
certainXprefer_dark	-1.434	0.720	-1.993
uncertainXprefer_dark	-0.187	0.717	-0.261
not certain	-0.402	0.497	-0.808
prefer_milk	0.625	0.480	1.302
uncertain	-0.113	0.524	-0.215
second bid	0.368	0.049	7.575
familiar	0.261	0.166	1.577

All the evidence thus speaks strongly in support of Hypothesis 2 even after other influences on bidding are controlled for.

The test of Hypothesis 3 concerning the moderating influence of group size is complicated by the apparent main effect of group size. As seen from the control condition, bidders seem to bid lower when there are more of them, an effect not predicted by the current independent private-value theory. One theory that would explain such an effect is a common-value theory: if the chocolates in fact have a common value element, bids should decrease with the number of bidders in order to avoid the winner's curse. Given this unexpected effect of group-size, it is not possible to say more about the influence of group-size on bargain-hunting. Only weak support for Hypothesis 3 follows from the fact that in the small-group (four bidders) condition, the difference between control and future-aware bids is 1.17 (one-tailed t-test p=0.02), whereas in the large group, the difference is 0.19 (p=0.36). However, with such a difference in the control-condition bids, Hypothesis 3 is not confirmed.

Conclusions from Experiment 1

The experiment found evidence for two theoretical predictions and remains inconclusive about the third. When similar but not identical substitutes are sold in a sequence of auctions, the amount of bidders' knowledge of future auctions significantly and often strongly influences their bidding strategies. Bidding is not only influenced by bidder awareness of a future auction, but also by the precision of information about the object to be auctioned in the future auction. As predicted by theory, bidders aware of a future auction bid less than unaware bidders. More precise knowledge of the object auctioned in the future makes the bidders polarize their bidding strategy towards the preferred object. The number of bidders participating in the auction does not seem to moderate the effects of a future auction on current bids. To summarize, the future opportunity to buy substitute products affects bidding in a way consistent with the theory, but the bidders do not seem to correctly take into account the intensity of the future competition as signified by the number of bidders per auction.

Limitations of Experiment 1

While Experiment 1 goes directly to the heart of the future-knowledge effect, the experimental design has several weaknesses and limitations. First, the fact that the control-group subjects have to be surprised by the second auction does not allow repeated measurements. Therefore, it is hard to give the subjects experience with the situation and ensure that they understand how a second-price auction works. It seems that at least a few subjects misunderstood the second-price auction as "I don't have to pay what I bid, so I can bid a lot to make sure I win", and ended up bidding very high. Experience with several rounds of bidding would most likely correct this misunderstanding. Second, the unit-demand assumption in the theory is not explicitly imposed on the subjects. The design therefore measures bargain-hunting in a realistic context of imperfect substitutes, but at the cost of being unfair to the theory because the predicted effect is likely smaller than in the case of perfect substitutes. It would be interesting to run another condition, in which the winner of the first box would not be allowed to bid again. Finally, the third and most important weakness of the design is that there is no control of individual valuations. Since the theory is ultimately about the mapping from valuations to bids, having only noisy and binary information about relative preference for the second object is not sufficient for detailed tight testing of the theory's predictions. For example, it is impossible to test for the predicted concavity of the expected surplus function, and it is impossible to examine the fit of the predicted bidding function to the actual bidding data. All three of the above limitations are addressed in Experiment 2 which I discuss next.

Experiment 2: Hypothetical product, individual valuation controlled

The theory of strategic bidding in two sequential auctions makes a very specific prediction about the relationship between the first-period bid and the private values of the objects being sold: $b_1(v_1,v_2) = \max[0,v_1 - S(v_2| b_1(v_1,v_2))]$. In particular, the first-period bid should be decreasing and concave in the second object's private value. Since under the distributional assumptions of the theory S increases in both arguments, there exists some S^* such that: $b_1(v_1,v_2) = \max[0,v_1 - S^*(v_2,v_1)]$, and the exact S^* can be numerically approximated. Please see Figure 1.2 for an illustration of the bidding function.

To test this theory tightly, it is necessary to know bidders' private values and, more fundamentally, to know that the bidders have fixed private values in the first place. One way to know both is to induce known values in a laboratory setting. Private values can be induced in the sense that the experimenter sells hypothetical and otherwise worthless objects in an auction, and the bidders are paid according to the surpluses they earn.

I sold two hypothetical products called "widgets" under two conditions: the two widgets were either substitutes or unrelated goods. All details of the game mimicked the assumptions of the theory:

- the two auctions were held sequentially, the first finishing before the start of the second
- the auctions had the second-price sealed-bid format
- both widget-values were privately known before the start of the first auction
- the bidders were symmetric in their valuation distribution (iid Uniform[4,9])
- the bidders were symmetric in their knowledge of the valuation distribution
- the number of bidders in each auction was fixed and known to be four.

Moreover, in the treatment condition, the winner of the first auction did not bid on the second object, so the unit-demand substitution assumption was implemented.

The theory predicts a measurable difference in bidding on the first widget. Given the parameters of the game, the bidders in the control condition have a dominant strategy to bid their valuation v_1, whereas the bidders in the treatment condition should bid exactly $v_1-S^*(v_2,v_1)$. No zero bids are expected under these parameter settings, because the maximum bid-decrement is bounded from above by the "myopic" M(1) function as illustrated in Figure 1.3, and for the assumed parameter values, M(1)=5/3, which is less than the minimum v_1 of 4.

To give subjects experience with the bidding environment and to study learning, the two-period sequential auction was repeated twenty times with a random secret re-shuffling of groups between iterations to prevent any repeated-game issues from polluting the results. The subjects were not informed at all about their opponent's identities, they only knew their number (three), and that the assignments to groups were random. I will now outline the experimental procedure.

Experiment design

The experiment was held in a computer lab on a university campus, and the subjects were mostly university students and staff. The subjects were asked to focus solely on the task at hand, and to refrain from using the computers for any other purpose. Each session followed seven steps:

1) Detailed written directions shown in Appendix 1.5 were given to the subjects in hardcopy along with two pages of screen-shots to familiarize the subjects with the bidding environment. The underlined sentence appeared only in the directions of subjects in the treatment condition, with the computer bidding environment adjusted accordingly. To capture all data for a potential study of learning and equilibrium-evolution, there were no practice rounds. Instead, all auctions counted towards the final payoff.

2) After giving the subjects about ten minutes to read the directions, all questions were publicly answered, and the experimenter went in detail over the two novel abstract

concepts used in the experiments, namely values and second-price auctions. The same simple concrete example of a second-price auction was always used to demonstrate how the winner and the price were determined.

3) The subjects had to correctly answer two diagnostic four-choice multiple-choice questions aimed at probing their understanding of the relationship between values and payoffs, and their understanding of price-determination in second-price auctions. The first question provided a randomly-drawn scenario with three concrete bids and asked the subject to correctly determine the winner and the price paid. The second question provided a concrete example of a bidder with a fixed valuation paying a price for the product, and asked the subject to correctly identify the earnings of that bidder in terms of surplus.

4) After all subjects answered the diagnostic questions correctly, they proceeded to the first bidding round, in which they bid on the first two widgets. The experimenter talked about every screen that appeared in the process to further explain what the subjects were looking at. The only feedback provided after every auction was whether or not the subject had just won a widget, and a playback of the widget's value and the price. A summary of total earnings so far was also displayed.

5) For the next 19 rounds, the experimenter did not interfere with the experiment unless there were questions or technical problems with the computers.

6) In the end of the experiment, the subjects were asked to summarize their bidding strategy as accurately as possible.

7) The subjects were then paid their earnings in cash and dismissed.

Results

56 subjects participated in the experiment, 24 in the control condition, 32 in the treatment condition. Because of technical difficulties, 16 subjects in the treatment condition only completed 18 rounds of trading. The treatment condition actually had two sub-condition because the first group of twenty subjects was not allowed by the system to bid over the maximum possible value, i.e. $9. This restriction was removed for all other groups, but it does not seem to affect the results. The constraint was binding for only a few subjects in a few rounds, and those observations are thus effectively censored. The mean bid in this sub-condition is actually slightly higher than the mean bid in the unconstrained control condition, so pooling the data across the two sub-conditions does not qualitatively affect the conclusions. I now discuss the main findings.

From past results (reference here), we know that people have a hard time doing the right thing in second-price auctions. Even with written and verbal explanations and examples how negative earnings arise and after some experience, many people tend to overbid. For the purposes of a simple first-pass analysis, I therefore discard the first ten rounds of bidding as practice for the subjects, and focus on the rounds 11-20. The descriptive statistics suggest that bargain-hunting occurs:

Table 1.7: Descriptive statistics, bid on first object in rounds 11-20

	TREATMENT	CONTROL
Number of first bids	294	240
Mean first bid	6.46 (0.11)	6.80 (0.10)
Median first bid	6.58	6.76
Mean (bid-value)	-0.07 (0.065)	0.29 (0.034)
Percent (bid<value)	30	4

These overall statistics confirm the general prediction that the treatment should reduce the bids on first widgets, although the standard errors are probably too small given that multiple observations per person are used to compute the means. The question remains, however, whether the bidders follow the optimal equilibrium strategy. The scatter-plots in Figure 1.5 show that the story is quite complicated. The plot below shows the bids on the first widget normalized by the private values of the first widgets (omitting a few very high outliers, all in the control condition). The red concave curve in the treatment condition depicts the approximate average optimal strategy.

Figure 1.5: First Bid as a Function of v_2, rounds 11-20, all subjects and observations

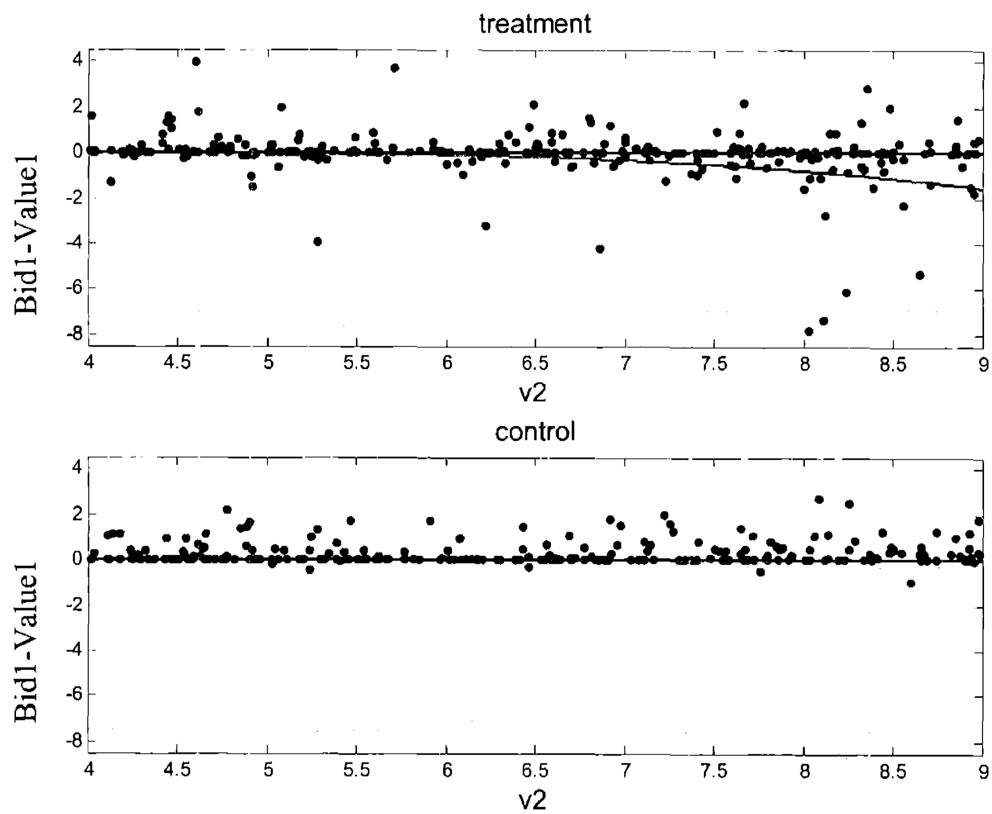

It is clear that the detailed prediction does not hold up in general. Instead, several distinct groups of people emerge from the data:

1. Some subjects persist in bidding their valuation even in the treatment condition. From reading the stated strategy-descriptions provided by the subjects and matching them with subject's majors, it can be guessed that economics training can lead to this strategy – the familiar dominant-strategy result is just too easy to latch onto.

2. Some subjects understand the need to bid lower in expectation of future profits, but end up shading their bids too much. Insufficient shading is also likely, although indistinguishable from just bidding valuation.
3. Some subjects just do not understand the game enough, and they persist in bidding over their valuations.
4. Finally, some subjects bid roughly along the line predicted by theory.

To study the fit of the theoretical model statistically, I can estimate the fit of the theoretical optimal equilibrium strategy by regressing the bids on three variables:

- *constant*: reflects the general shading of the bids irrespective of future or current values
- v_1 : The value of the first widget, expected coefficient is unity
- $S^*(v_2,v_1)$: the numerically approximated equilibrium expected surplus function, the expected coefficient is negative unity

In summary, the ideal treatment-condition data would have a perfect fit and coefficients [0,1,-1]. Before estimation, I feel it is important to ignore some of the subjects who obviously do not understand the game even with practice. One way to screen subjects fairly across the experimental condition is to use their bids on the second widgets. Since the second-period optimal strategy is the same in both conditions, screening on the basis of the second-period bids does not apriori favor one condition over another. Keeping only the subjects who never overbid more than a dollar on the second widget after round ten, and thus eliminating 6 subjects in the treatment condition and 5 in the control condition, the regression results are as shown in Table 1.8.

Table 1.8: Fit of the theoretical model

	Treatment condition			Control condition		
		N=236	R^2=0.713		N=190	R^2=0.973
Variable:	coefficient	st. error	t-value	coefficient	st. error	t-value
Constant	-0.714	0.300	-2.377	0.083	0.083	1.001
v_1	1.101	0.046	23.860	1.003	0.013	78.264
$S^*(v_2,v_1)$	-0.713	0.192	-3.707	0.011	0.049	0.233

From the two above regressions, it is clear that at least some of the bid-shading observed in the treatment condition from the overall descriptive statistics is constant with respect to v_2. An average bidder sheds 71 cents in the treatment condition regardless of the second widget's value, and the second widget's value then has about 70 percent of the impact than it should. Therefore, when v_2 is low, the bidders shed too much and vice versa. While this interpretation of the regression coefficients is reasonable, it does not represent reality well because the "average bidder" it describes does not exist. Instead, as discussed above, there seem to be different groups of bidders, each using a different strategy. It is likely that the parameter estimates are driven by a few very low outliers – bidders who bid zero in the first auction in anticipation of a profitable second auction. Once these seven outliers are removed (by only accepting observations, in which the difference between value and bid does not exceed $3), the following regression tells a different story:

Table 1.9: Fit of the theoretical model, seven outlying observations removed

	Treatment condition, (bids futher than $3 from value removed)		
		N=229	R^2=0.92
Variable:	coefficient	st. error	t-value
constant	-0.425	0.142	-2.995
v_1	1.065	0.022	49.002
$S^*(v_2,v_1)$	-0.355	0.091	-3.895

Both the intercept and the coefficient on the surplus function are almost halved in size. Most importantly, however, the R^2 statistic rises dramatically, and only 8 percent of the variance now remains unexplained by the model. Therefore, removing the few bidders who sometimes underbid drastically increases the fit of the model, and produces more interpretable coefficients. Note that this removal of outliers is equivalent to zooming the bid-plots to within $3 of valuation. Please see Figure 1.6 for the resulting look at the data.

Figure 1.6: Zooming in to within $3 of value

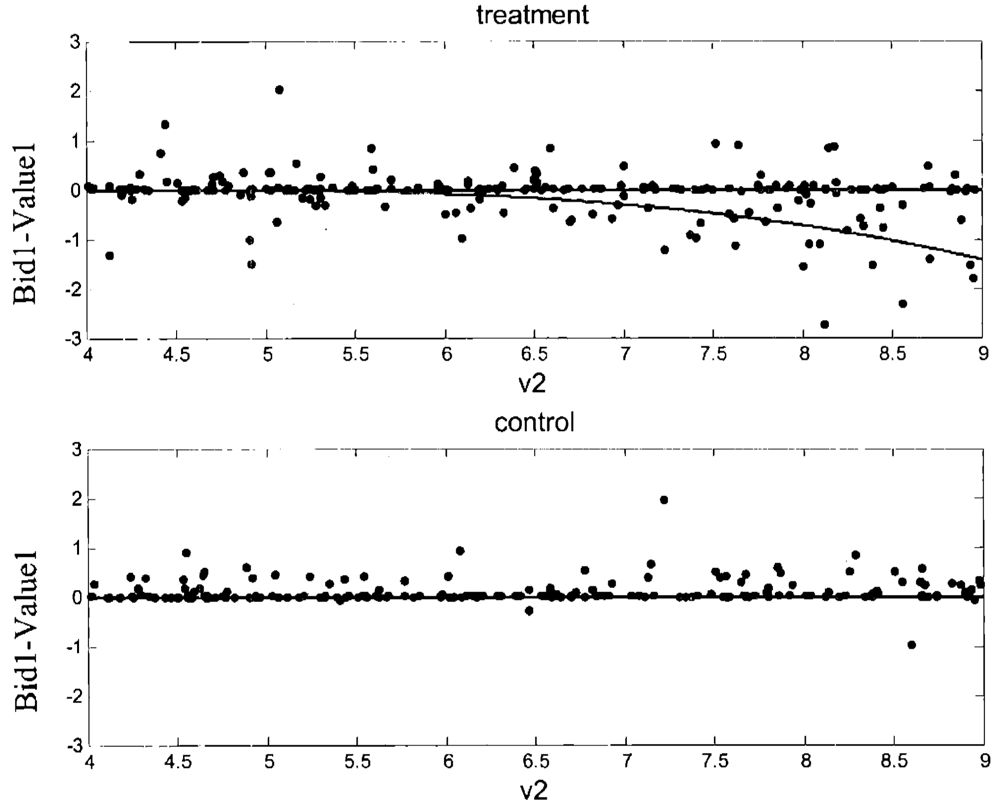

The analysis so far has considered the data at the bid level, in effect assuming that all bidders use similar strategies. But the preliminary look at the data and the written strategy descriptions provided by the subjects suggest heterogeneity in strategies. Therefore an analysis at the bidder level may be more appropriate. For the purposes of this analysis, we can "zoom back out" and include all 32 bidders in the treatment condition. The analysis has to remain preliminary because of the low number of observations.

One way to summarize each bidder's strategy with just two numbers relevant to the model is to run individual-specific regressions of the bids' deviation from values (bid-v_1) on a constant and the $S^*(v_2,v_1)$ function. The coefficient on the constant is a measure of a general over/under bidding, while the coefficient on $S^*(v_2,v_1)$ is a measure of bid-shading normalized to the bid-shading predicted from theory. Since there are ten observations per bidder in the advanced rounds of the game, the coefficients are not estimated with great precision, but they nevertheless provide an approximate summary of each bidder's strategy. Including the constant allows the regression to capture bid-shading according to v_2 even in bidders who generally bid too much as well as constant bid-shading in bidders who ignore v_2.

The results of the bidder-level analysis are best summarized graphically by plotting each bidder as a point in the (constant,S^*) parameter-space. In Figure 1.7, the large red dot marks the location of the optimal strategy, and the two outlying bidders are shown as close as possible to their actual location. Italics are used to interpret regions of the space as strategies. Finally, the parameters of two bidders who bid exactly their value all the time have been jittered slightly to separate the individuals visually.

Figure 1.7: Individual Bidders in (c,S^*) strategy-space

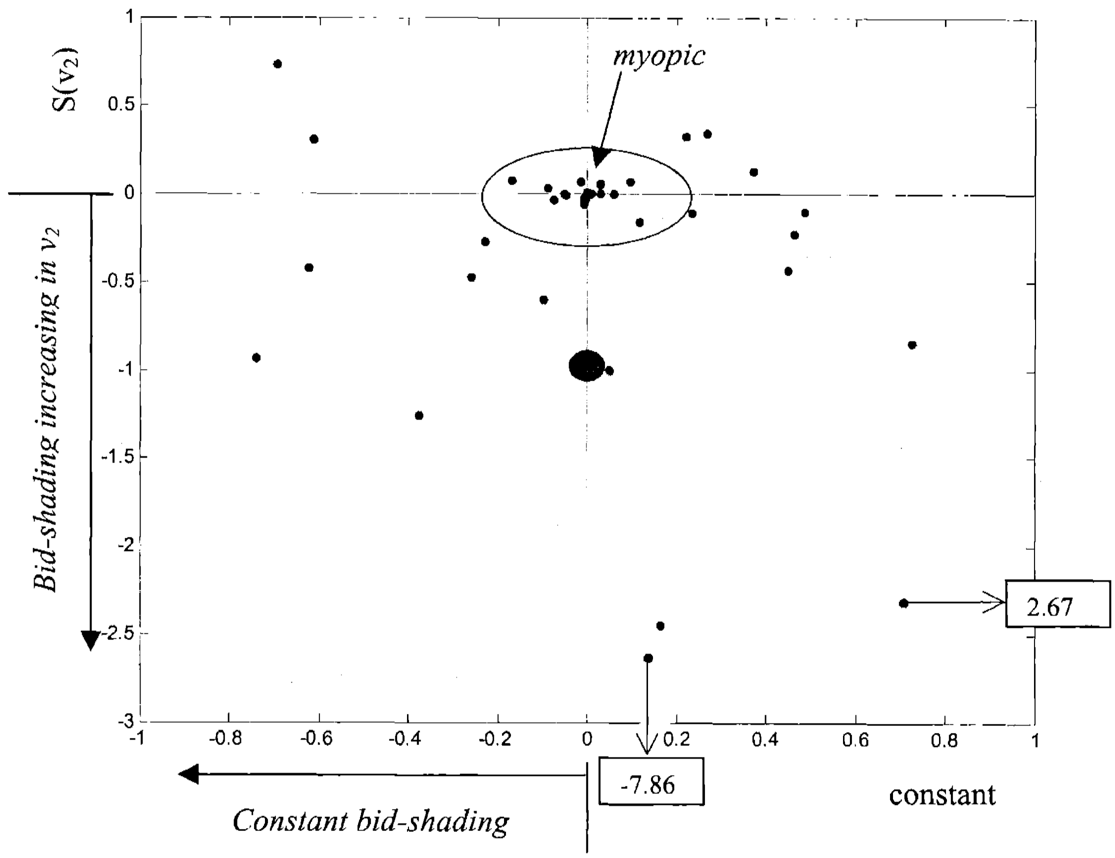

One interpretation of the results is the following: 12 bidders bids approximately their valuation, essentially behaving myopically. 3 bidders overbid and increase their bids in the second-object's value, a hard-to-understand strategy. The remaining 17 bidders engage in some form of bid-shading. Looking along the y-axis, we can see that of the 17 bidders who shade their bids, all but two decrease their bids as a function of v_2, but most do not decrease the bids enough relatively to the optimal equilibrium strategy. There are three bidders, however, with S^* coefficients –2.2, -2.4, and –6.3 respectively, who decrease their bids too much, actually bidding zero whenever their v_2 is high enough. Looking along the x-axis, we can see that there is widespread over-bidding, because nearly half of the bidders have a positive constant coefficient. Therefore, bids decrease in v_2 in most bidders, but only about half of the bidders actually submit some bids less than their valuation. The simplest summary of the heterogeneity in strategies is that about a third of the bidders use the myopic strategy of bidding valuation, a third of the bidders shade their bids below valuation and bid along a decreasing function of their second-object valuation, and a third of the bidders use strange strategies characterized by overbidding while most often still decreasing their bids as a function of their second-object valuations.

Conclusions from Experiment 2

Bids on the first of two perfect substitutes are lower than on the first of two unrelated goods, so bargain-hunting occurs in the experiment. Moreover, the value of the second substitute product reduces bids on the first product, but on average not as much as predicted by theory. Overall statistics can be deceptive because the experiment also reveals heterogeneity in strategies used by different bidders. Instead of a description of an average bidder, the data is better captured by a story, in which bidders cluster into several groups, each group using a distinct strategy. Some shade their bids too much, but only when the value of the second substitute is high, effectively skipping the first auction. Others decrease their bids slightly regardless of the particular value of the second product. Yet another group fails to decrease their bids at all, behaving myopically. By round ten, only about one third of the bidders are reasonably near the optimal forward-looking strategy predicted by theory, and only about half of the bidders sometimes bargain-hunt and submit bids lower than their valuations. It remains an open question whether with sufficient experience, all bidders convert to the correct strategy.

Econometric Evidence: Analysis of eBay data

The two-period model with a fixed pool of bidders discussed in the theory section is not immediately extendable to the eBay situation, where hundreds of substitute items are auctioned off every day, and the bidder pool is uncertain and constantly replenished. As shown on the three-period example, building models of many auctions is complicated if bidders try to interpret past prices and patterns of winning and losing. In a world with a fixed pool of bidders, such learning can be beneficial to the bidders, but the uncertainty and fluctuation of the bidder pool on eBay most likely reduce the informativeness of past prices about future competition. It is therefore without much loss of realism to abstract from the informational dependencies between past outcomes and future competition, and assume that bidders have no memory. Also, the effect of highest competitive bid on the expected surplus function is quite small in the two-period model, so it should be a reasonable simplification to ignore it, and assume that current valuations do not enter the calculation of the expected-surplus from future auctions. These two simplifying assumptions allow a tractable model of forward-looking bidders facing a very long stream of auctions to inform econometric analysis of bargain-hunting on eBay. I now outline the model.

Consider an infinite stream of auctions for substitute goods. To capture the idea that bidders always know the next few products coming up for sale, assume that at every time-period, each bidder knows his/her private single-item valuations $(v_0, v_1, \ldots, v_K)$ of the current and the next K products for sale, where the indexing of the products is relative to the current time period (v_0 is the valuation of the current product, v_1 is the valuation of the next product, etc…). To abstract from bidder learning, assume that the bidders do not remember past prices or valuations. In this sense the bidders are purely forward-looking. The valuations are stochastically equivalent and independent across bidders, each bidder wants at most one item, bidders are risk-neutral, and they discount the future by a factor $\delta<1$. For simplicity, assume that the number of bidders is unknown and varies randomly around a fixed mean, so there is no time-trend of the number of bidders. Then, it is reasonable to assume there is a steady-state distribution of competitive prices G(p) arising from actions of other bidders and a steady-state future surplus function S(**v**). For convenience, assume G has [0,1] support. Then, S(v) satisfies the following Bellman equation:

$$S(v_0, v_1, \dots, v_K) = \max_{\beta \geq 0} \left\{ \int_0^\beta (v_0 - p) dG(p) + \delta E_w [S(v_1, \dots, v_K, w)](1 - G(\beta)) \right\} \tag{1.7}$$

where β is the current bid and w is the still unknown valuation of the item K+1 periods ahead. The current-period maximization problem is solved by[12]:

$$\beta = \max\{v_0 - \delta E_w [S(v_1, \dots, v_K, w)], 0\} = \max[v_0 - \delta ES(v_1, \dots, v_K), 0] \tag{1.8}$$

where the ES notation stands for "expected surplus".

Therefore, the bidder only collects incremental surplus in periods with high-enough valuations[13]. If valuation is less than ES, the surplus must be less as well, so it is better to bid zero and wait.

Note that the steady-state optimal bidding strategy shown in (1.8) resembles the two-period strategy in equation (1.2) in that the bidder bids valuation less the expected surplus from the future auctions conditional on the known future valuations whenever that is a positive quantity. The ES function links the known information about future values to an assessment of future desirability, and it is also a measure of the opportunity cost of winning the current item. Obtaining a closed-form solution is difficult, but it is possible to show that the function is non-decreasing in all its arguments, future is less important than today, and the more distant is the future object, the less impact does its valuation have on the expected surplus:

---

[12] The result follows from first-order conditions. The problem is concave at the solution as long as g(β)>0.

[13] Note that the bidding function suggests we have to allow F(p) with a spike at zero. Would it ever be profitable to bid a small ε to possibly get the good for free? Since this concern is only relevant whenever v<ES, the expected profit would be Pr(p=0)v which is less than ES, and the bidder is thus better off with biding zero and waiting. The tacit assumption is that only strictly positive bids can ever win the auction.

Proposition 1.4: The steady-state equilibrium of the K-period look-ahead game is characterized by an expected-surplus function $ES(v_1,...,v_j,...,v_K)$ with the following properties: ES is non-decreasing in all arguments, partial derivatives of ES are all smaller than unity, and they on average decrease monotonically with j. (proof in Appendix 1.6).

Note that as shown in the theory section and demonstrated in the experiments, the assumption about bidders' knowledge of future auctions plays a key role in the bidding strategy. Manipulating this assumption produces two alternative models: the "uncertain" model of Engelbrech-Wiggans (94) and the benchmark myopic model. The uncertain model[14], involves only the current valuation and shades down the bid according to net present surplus of the future – a constant in terms of the yet-unknown future valuations. The myopic model does not shade down the bid at all. These differences provide the basis for an econometric test that will distinguish among the three models. The differences among the three candidate models are illustrated in the Table 1.10 below:

Table 1.10: Predictions of different models of future knowledge

Model of future knowledge in sequential auctions	***What about the future influences today's bid?***	
	General properties	Specific future items
Certain (future items known)	✓	✓
Uncertain (future items uncertain)	✓	
Myopic (no regard for the future)		

[14] Really a special 0-period look-ahead case of the K-period model.

Data

EBay was kind to provide data on bidding in the MP3 player category. The data spans several months in the spring of 2001, and it provides rich information about each auction, including the starting, reserve, buy-it-now, and final prices, text description of the item sold, and the seller's marketing features like photos or bold font. In addition, the data gives the researcher an insider's view of every bid made in all of the auctions. Therefore, we know all the proxy bids made including the winning bid, which remains undisclosed in reality (please see www.ebay.com for a definition of proxy bids and a thorough description of the bidding rules). Much like in a supermarket panel dataset, a numerical id of each person is listed with each bid and allows us to track the same bidder across different auctions. But unlike in a supermarket panel dataset, the items for sale are not pre-classified and have to be inferred from user-entered text descriptions.

The choice of product-category was driven by two major considerations: First, MP3 players are relatively expensive gadgets likely to be subject to the unit demand assumption used throughout the theory section. Second, the MP3 player product space is dominated by only a handful of brands, each producing just a few products, so it is relatively easy to discern which product is being auctioned from the text description available in the data. To achieve this product-identification, I use a word-matching algorithm that looks for brand and model names and finds the product that best fits the text description. Whenever the best fit is poor, the item for sale remains unclassified. About 15 percent of all auction listings remain unclassified, either because their description is insufficient for identification ("new cool mp3 player for sale") or because they do not belong to the product category at all ("napster t-shirt" or "128mb sandisk memory card"). For the purposes of estimation, I distinguish among the 36 most listed products that make up 95 percent of the listings. The remaining few rare products and all the unclassified items are all classified together as "Other".

The data provided by eBay comes with one more imperfection – only the ending date of each auction is reported, not the ending minute. I therefore impute actual time using all available information. Unlike the product-identification described above, this imperfection is accidental and not a structural property of online auction data, so I relegate the details of my imputation procedure to the Appendix 1.7. It suffices to note that I was able to guess exact times roughly within an hour of the true times and, more importantly for the analysis below, the imputed times capture the order in which the auctions ended almost perfectly because the order information is contained in the auction identification numbers.

Before using the data for model estimation, I carried out several data-cleaning steps. First, I eliminated all auctions with missing data, bid retractions or cancellations, as well as all private and featured auctions. The featured auctions were eliminated because there were very few of them, not enough for a control within the model. Then, I eliminated the few auctions in which I was unable to re-create the listed final price from the bidding history, an evidence of missing data. Finally, I eliminated all auctions without any bids. In the resulting dataset, there were 8435 auctions, and participating in these auctions were 25785 bidders, 57.67 percent of whom just bid in one auction. Among the single-auction bidders, 25.18 percent are eventual winners, while among the multi-auction bidders, 35.44 percent are eventual winners. These numbers suggest a relatively thin market with about three bidders per item on average. Note that since all sequential-auction effects are larger in thinner markets, measurement of the effects should be possible here. Of the 3868 winners who participated in multiple auctions, 84.31 percent won exactly once, 12.93 percent won exactly twice, and 2.77 percent won more than twice. The very high proportion of single wins is evidence for the unit-demand assumption. Throughout, I only consider the bidders who won at most twice; the remaining bidders are probably institutional buyers likely facing different incentives. The average number of auctions a bidder participates in is 2.50. Note that this figure defines participation as submission of a bid. Since a bid can only be submitted if it exceeds the highest bid at the moment, the actual participation figures are likely to be much higher, perhaps by an order. To complete the data cleaning, I eliminate outlier bids (top and bottom 5% of all bids on each product). To align the data the sealed-bid theory, I consider only each bidder's highest bid in each auction. So whenever a bidder bids multiple times in an auction, I interpret his highest (and by construction final) bid as the bid they would have submitted in a sealed-bid situation. From now on, this highest bid is what I mean by "the" person's bid in a particular auction.

Finally, I consider on "serious" bids by eliminating bids that occur earlier than the last day (24 hours) of the auction (roughly 45% of bids). Considering only late bids brings the sealed-bid theory closer to eBay's open ascending bid reality, in effect modeling only the eBay endgame. The present theory does not model strategic early bids, and thus does not capture the behavior of bidders who are not present during the endgame when most of the winner-determining bidding happens. To avoid autocorrelation and other dependency issues involved in the same person's bids, I only draw one bid per person at random for the estimation dataset. This data-cleaning procedure produces 14181 observations, 3384 of which are bids by bidders who bid in three auctions or more.

Econometric Model

To measure the effect of future knowledge on bidding while distinguishing among the three candidate assumptions about future knowledge, I estimate the following conceptual model of bid of person i in auction a:

$$\text{bid}_{ia} = \alpha v_{ia} + \beta v_{i,\text{future}} + \gamma v_{\text{future}} + \theta \text{controls}_{ia} + \varepsilon_{ia} \qquad (1.9a)$$

where v_{ix} is the valuation of person i of x, v_{future} is a measure of the general population value of the future, and the $controls_{ia}$ variables include product-specific dummies along with several other control variables discussed below. Finally, ε_{ia} is a mean-zero error. It is easy to see that the hypotheses corresponding to the candidate assumptions are:

myopic $\Rightarrow$ $\beta=\gamma=0$, ***uncertain*** $\Rightarrow$ $\beta=0$ & $\gamma<0$, ***certain*** $\Rightarrow$ $\beta<0$ & $\gamma<0$

To really specify the model, I need to define $v_{i,future}$. The general class of certain models calls for some aggregation of bidder i's valuations of the near-future objects. In particular, a simple average of the future valuations would be reasonable. The K-period look-ahead model provides further structure in addition to the general certain model. In particular, if the K-period look-ahead model holds, then the econometric model should be specified as:

$$\text{bid}_{i,a} = \alpha v_{i,a} + \sum_{J=1}^{K} \beta_J v_{i,a+J} + \gamma v_{future} + \theta \text{controls}_{ia} + \varepsilon_{ia} \qquad (1.9b)$$

and the following should be true for all J: $|\alpha|>|\beta_J|$, $|\beta_J|>|\beta_{J+1}|$, and $\beta_J<0$. I will consider both the (1.9a) specification (using a simple average of future valuations) and the (1.9b) specification, but before discussing the detailed specification of the model, it is necessary to discuss where the valuation variables come from.

Constructing individualized valuation variables

The individual valuation variables v_{ix} are central to the theory, and since they are not directly observed, they need to be constructed from information available in the data. To increase the validity of the results, I construct several different valuation variables and make sure that the qualitative conclusions are not sensitive to the choice of a construct.

Valuation Construct 1: Current bid as a signal of preference

Since bidding implies commitment to eventual purchase, the choice of a particular item to bid on reveals a relative preference for that item. Therefore, assume the item sold in auction a is bidder i's "favorite " item. Now, bidder i's favorite item being sold in the near future (next few auctions) increases the relative value of the future to that particular bidder. Using just a binary indicator variable for the favorite product being sold in the near future thus allows estimation of the following constrained conceptual model:

$$\text{bid}_{ia} = \beta \mathbf{1}[v_{i,future} - v_{ia} > c] + \gamma v_{future} + \theta \text{controls}_{ia} + \varepsilon_{ia} \qquad (1.10)$$

The main benefit of this simple variable is that it can be constructed even for bidders observed only once. Therefore, the full sample of 14181 bidders can be used in estimation.

Valuation construct 2: Past bids as signals of preference and category valuation

Following the revealed-preference logic behind construct 1, bidders who submit bids in several different auctions provide more information about their preferences for products. In particular, selecting one of the bids as the "current" bid allows the modeler to look into the bidder's "past" for valuation-relevant variables and construct both v_{ia} and $v_{i,future}$ to estimate the unconstrained model (1.9a). This gain in information comes at a cost, however, because bidders who only bid once cannot be used. To increase precision of the constructs, I only use the data from the 3384 bidders who bid in at least three auctions. Note that the bidder's "past" also provides information about bid amounts for several products, but this information is likely not exogenous enough to be used as an explanatory variable of the current bid amount. Therefore, I focus on constructing relative valuation variables from the relative preference information contained in each bidder's bidding history.

Figure 1.8: Constructing personal valuations from personal bidding history

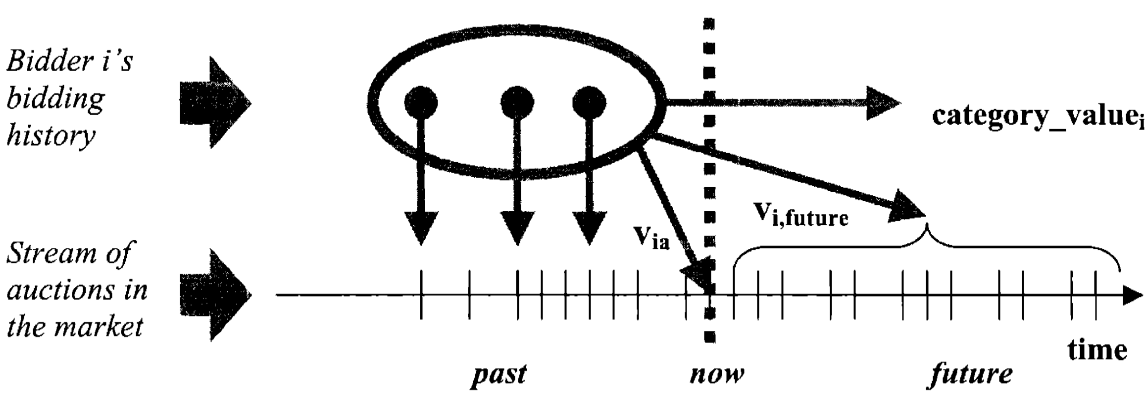

The relative product-preference information consists of a vector of "choice" shares for each bidder. To construct valuations theoretically consistent with the observed shares, I use the logit model. Assuming the logit model and fixing one of the valuations arbitrarily at unity, there is a 1-1 mapping between the shares and the product-specific valuations that would produce:

$$\text{share}_{ip} = \frac{\exp(v_{ip})}{\sum_{k \in products} \exp(v_{ik})} \text{ and } v_{i1} = 1 \Rightarrow v_{ip} = 1 + \log\left(\frac{\text{share}_{ip}}{\text{share}_{i1}}\right) \tag{1.11}$$

Since each bidder's shares are just crude small-sample estimates of the true shares of preferences, a more accurate individual preference variable can be obtained by incorporating the population's preferences as well. Moreover, $share_{i1}$ is bound to often be zero for even the most popular product 1, so the above formula (1.11) is not always defined. To improve the share estimates and make sure formula (1.11) is always defined, I compute a convex combination of each bidder's share and the population share before applying formula (1.11). To give more weight to the bidder's shares whenever they are more precise, I use the inverse of the size of the bidder's history as the population share's weight as shown in equation (1.12). Note that this weighting scheme is not arbitrary, the statistical precision of the share estimate is exactly proportional to this weight.

$$\text{share}_{ip} = (1-\lambda)\frac{1}{N_i} \sum_{k \in past(i)}^{N_i} 1(\text{choice}_{ik} = p) + \lambda \text{share}_p \quad \text{where } \lambda = \frac{1}{N_i} \tag{1.12}$$

The resulting valuation variables are well suited for a linear model because they are approximately symmetrically distributed and they do not have severe outliers.
Besides delivering a measure of each individual's relative valuations for all products, the bidding history also provides information about the individual's overall willingness to pay for MP3 players. To capture this information, I construct the $category\_value_{ia}$ variable by calculating an average bid as a proportion of the final price whenever the bid is less than the final price:

$$\text{category\_value}_i = \frac{1}{P} \sum_{i,past=1}^{P} \frac{\min(bid_{i,past}, price_{i,past})}{price_{i,past}} \tag{1.13}$$

The resulting variable is a number between 0 and 1 monotonic in the overall willingness to pay. Another way to interpret the variable would be that it captures individual-specific bargain-hunting propensity, thus providing an important control in a model that's explaining bargain-hunting as rational forward-looking.

So far, I considered preferences over products. I also considered preferences over brands and other product attributes like memory capacity, and the qualitative conclusions reported below did not change. All these variations on the constructs provide the same qualitative conclusions, so I report only results for the following variable settings: "future" is next 10 auctions (so K=10), preferences are over products, and they are calculated by equation (1.11) applied to adjusted shares calculated by equation (1.12).

Control variables

Besides the constant and 37 product-specific dummies, I use the following control variables:

- *future_freq*: Per-hour frequency of auction endings implied by the next 10 auctions. (captures product-independent desirability of the future, minimum 0.8, mean 5.9, median 5.2 and maximum 30)
- *bold, gallery and photo* : marketing mix variables at seller's disposal (binary)
- *timeleft*: Time left to the end of auction (in minutes, by construction decreases bids)
- *buyitnowprice*: "Buy-it-now" price, set to overall mean when option being available. Sellers can use this option as a spot-market price available before bidding begins, an incentive to set it relatively high. It can be used as an anchor. As such, it should increase bids.

Estimation issues: sample selection

One potential problem with estimation of equation (1.9) is that on eBay, a bid can only be submitted if it exceeds the highest bid at the moment. Therefore, the dataset contains relatively more high bids and relatively fewer low bids than a random sample of willingness-to-bid. Since personal heterogeneity in valuations is likely the main driver of bids, the dataset thus observes more high-valuation bidders than low-valuation bidders. A correction for this potential sample-selection bias is subject for future research. One way this sample-selection might bias inference from standard OLS would be to observe only bidders who deflate their bids relatively less when future is desirable. Therefore, barring other possible effects, OLS-based inference about β and γ should be conservative, and a potential finding of significantly negative parameter-estimates should be robust to the sample selection.

Estimation Results based on Ordinary Least Squares

In all the regressions below, observations correspond to individual bidders, and the dependent variable is a single randomly selected bid as described above in the "Data" section. Throughout the result-discussion, I suppress the constant and 36 product dummies for parsimony. I first discuss the results obtained from estimating the restricted model (1.9a) on the complete sample of bidders, please see Table 1.11.

Table 1.11: Full sample of bidders, restricted (1.9a) model

Variable	Coefficient	Error	t-value
log(mins_left)	-3.524	0.135	-26.169
photo	-2.378	0.753	-3.159
bold	5.457	1.486	3.673
gallery	0.221	1.140	0.194
buyitnowpri	0.322	0.007	44.502
future_freq	-0.579	0.084	-6.892
future_favorite	-3.668	0.752	-4.876
N=14181, R^2=0.74			

All variables have the expected signs, and both the general (*future_freq*) and the specific (*future_favorite*) future-focused variables are significant and sizeable when compared to the average bid in the dataset of about $100 (standard deviation about $70). The interpretation of the respective coefficients is straightforward: for every additional auction ending within the next hour, bids decrease by about $0.58, while the favorite product being offered at least once within the next 10 auctions decreases current bids by about $3.67. The same regression run only on the 3384 multi-auction bidders who bid in three auctions or more produces similar coefficient estimates and does not change the qualitative conclusions. Therefore, I conclude that the multi-auction bidders are representative of the whole population as far as forward-looking is concerned.

Of the marketing instruments available to the seller, only bold listing seems to have a discernible positive effect, yielding a bid increase of about $5.46 on average. Considering the fact that the bold-listing fee was $4 at the time of the study, it seems like bold listings are a good deal, but the endogeneity in bold-listing allocation prevents a definite conclusion. The buy-it-now price has the expected anchoring effect, increasing the price by $1 leads to an increase in bids of about 32 cents. The future_favorite variable can be used for a crude mean-based test of the effect of the future: The mean bid in auction where the same product is available in the next ten auctions is $116.38. The mean bid in auctions where the same product is not available in the next ten auctions is $105.49. The difference is quite significant given the large sample-size. Decomposing the effect of future valuations into lead-specific indicators of favorite product's availability leads to a restriction of the (1.9b) model. The results are shown in Table 1.12 below.

Table 1.12: Full sample of bidders, restricted (1.9b) model

Variable	Coefficient	Error	t-value
log(mins_left)	-3.547	0.134	-26.389
photo	-2.304	0.753	-3.061
bold	5.480	1.485	3.691
gallery	0.005	1.138	0.005
buyitnowpri	0.321	0.007	44.452
future_freq	-0.518	0.084	-6.152
favplus1	-5.870	0.922	-6.365
favplus2	-3.352	1.022	-3.278
favplus3	-3.152	1.034	-3.048
favplus4	-1.040	1.011	-1.028
favplus5	-0.821	1.024	-0.802
favplus6	-1.594	1.036	-1.538
favplus7	-1.370	1.034	-1.325
favplus8	-0.760	1.059	-0.718
favplus9	-1.994	1.040	-1.918
favplus10	-0.119	1.076	-0.111
N=14181, R^2*=0.74*			

In agreement with the K-period look-ahead model (Proposition 4), the impact of the current (here assumed to be favorite) product being available J periods into the future declines with J both in magnitude and in significance. The trend is obvious, statistical testing is not necessary. Less surprisingly, all coefficients are negative, so future availability of the current product reduces the current bid.

While the full sample is most representative of the bidder population, it does not allow the modeler to construct individual relative-preference and category-valuation variables, thus making estimation of models (1.9a) and (1.9b) impossible. Focusing on the sub-sample of 3384 bidders who bid in at least three auctions, these variables explain a significant portion of the variance and allow more precise hypothesis testing. As in the case of construct 1, all significant variables have the expected signs and both future-focused variables are significant. Please see Table 1.13 below for the coefficient estimates.

Table 1.13: Bidders in at least three auctions, model (1.9a)

Variable	Coefficient	Error	t-value
log(minleft)	-2.846	0.238	-11.974
photo	-0.694	1.158	-0.599
bold	5.827	2.674	2.179
gallery	0.105	1.993	0.053
buyitnowpri	0.225	0.013	16.984
now_prod	1.884	0.342	5.513
future_prod	-2.067	0.578	-3.574
future_freq	-0.484	0.133	-3.651
category_value	39.338	2.691	14.619
N=3384, R^2*=0.81*			

The general future effect remains fairly large: for every additional auction ending within the hour, bids decrease by about 48 cents. To gauge the size of the specific future effect, note that the *future_prod* variable is roughly centered at zero with a standard deviation of 1.14. Therefore, increasing the average value of the next ten auctions by one standard deviation decreases the current bid by $2.35. Another way to interpret the size of the effect is the absolute value of the ratio of *future_prod* coefficient and *now_prod* coefficient. The result is impact of preference for specific future products relative to the impact of preference for the current product. In this sense, the specific future-effect is quite large. Decomposing the effect of future valuations into lead-specific future product valuations allows estimation of the (1.9b) model. The results are shown in Table 1.14 below. Consistent with the K-period look-ahead model, the impact of future valuations is negative and declining in both magnitude and significance. Also as predicted, the downward impact of each future valuation is smaller than the equivalent upward impact of the current valuation.

Table 1.14: Bidders in at least three auctions, model (1.9b)

Variable	Coefficient	Error	t-value
log(minleft)	-2.871	0.238	-12.062
photo	-0.705	1.161	-0.607
bold	5.953	2.685	2.218
gallery	0.093	1.997	0.047
buyitnowpri	0.225	0.013	16.965
now_prod	1.916	0.342	5.600
future_freq	-0.469	0.133	-3.537
category_value	39.391	2.693	14.629
vplus1	-0.708	0.281	-2.520
vplus2	-0.318	0.285	-1.117
vplus3	-0.615	0.287	-2.145
vplus4	0.049	0.287	0.172
vplus5	-0.117	0.291	-0.403
vplus6	-0.265	0.290	-0.915
vplus7	-0.276	0.295	-0.936
vplus8	-0.106	0.293	-0.362
vplus9	0.174	0.290	0.600
vplus10	0.092	0.295	0.313
N=3384, R^2=0.81			

Including the interactions between category value and the two relative preference variables would also yield an interesting model, but the correlations between the preference variables and their respective interactions with category value exceed 0.9, making an interpretation of the resulting coefficients impossible.

Conclusion of the empirical analysis of eBay data

Future opportunities matter to eBay bidders. All estimation results strongly reject the myopic model by finding large and significant bid reduction whenever the future is more desirable. Moreover, bidders consider not only the unconditional expected value of the future as the uncertain model would predict, but also condition their future assessment based on *specific products* available in the future. In particular, bid deflation is larger whenever the future has higher private value for the bidder. Therefore, the proposed theoretical model captures the forward-looking aspect of actual behavior in online auctions for MP3 players.

Discussion

This paper studies what happens when the role of an auction changes from selling unique objects at Sotheby's to driving entire markets for substitute goods on eBay. When participating in a sequence of auctions, rational bidders reduce their bids in anticipation of future auctions for substitute products. Moreover, the magnitude of the bid-reduction depends on the bidders' known private values of specific products sold in the future auctions. Actual bidders in lab experiments and in real online auctions seem to behave on average approximately as the rational bidders would except for two shortcomings: they do not correctly judge the impact of the number of competitors on expected surpluses, and there seems to be persistent heterogeneity in their individual strategies.

Since B2C and C2C auctions are being used to sell mass-produced consumer goods and the overall volume of auction-driven trading is rising, sequences of auctions for substitute products are a growing phenomenon on auction sites like eBay. In parallel, as B2B procurement auctions expand their scope, it is inevitable that more and more capacity-constrained suppliers are biddings sequentially on contracts of which they can only fulfill a few. Therefore, sequential strategic sequential bidding on substitutes should happen even in the economically larger B2B realm.

Besides the obvious normative implications for buyers, the findings reported here may have impact on seller strategies. There has been a dramatic increase in the use of auctions and one could ask whether the inter-auction competition analyzed in this paper will limit the extent to which auctions are used in the future. Understanding the supply side of the auction-driven markets is an interesting topic for future research that may shed light on this question of the scope of auctions. Throughout this paper, the seller was assumed exogenous, but allowing for strategic selling may both qualitatively and quantitatively change the bargain-hunting by bidders.

Finally, the findings have implications for third-party analysts of auction-driven markets in that they improve and clarify the interpretation of observed bids. In particular, bids in online auctions cannot be simply interpreted as unbiased signals of valuations, but must be considered in the context of other auctions for similar objects.

Appendices of Chapter 1

Appendix 1.1: eBay screen (MP3 players)

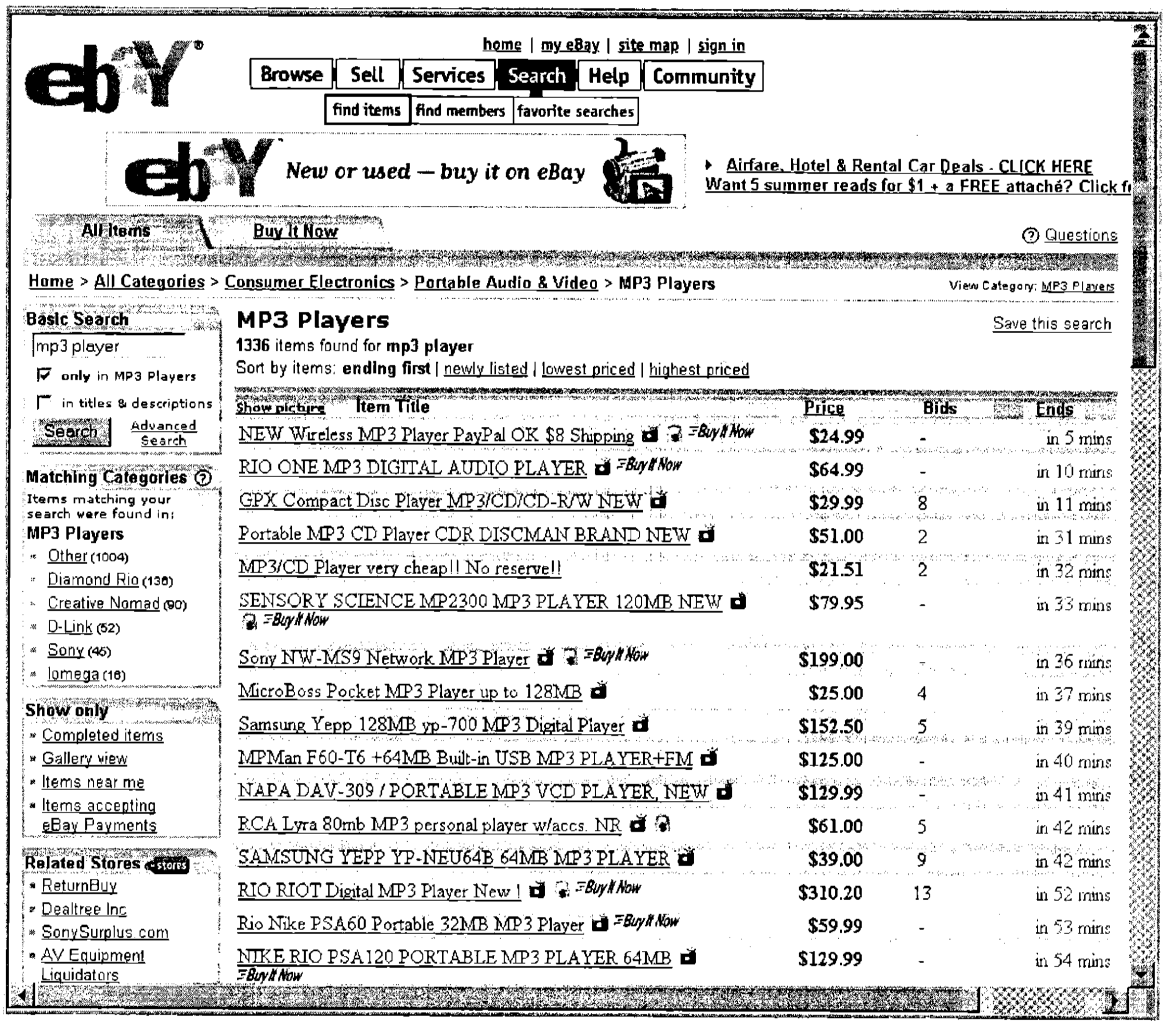

Appendix 1.2: Proofs of Theorems and Propositions

Proof of Theorem 1: Given any function S_K, let the mapping $T(S_K)$ map S_K to S_{K+1} such that $T[S_K(v_2)] = S_{K+1}(v_2, c) = \int_0^{v_2} G(p \mid c, S_K)dp$, where $G(p|c,S_K)$ is the distribution of second highest of the v_2's of first-period losers assuming first-period bidding proceeds conditional on S_K as in equation (1.3). I will use the Schauder Fixed-point theorem (Schauder 1930) to show that there exists S mapped to itself by T, hence an equilibrium S. The fixed-point theorem requires the following conditions for a mapping T to have a fixed point:

1) T: M→N where M is a nonempty convex subset of a Banach space X and N is a compact subset of M.
2) T continuous

X here is the space of all Borel-integrable functions from [0,1]x[0,1] into itself, and let M=N be the subset of X consisting of all continuous differentiable functions h with h(0,x)=0 and partial derivatives bounded weakly some fixed positive number B. This set N is trivially nonempty and convex.

Compactness is the non-trivial condition. By Arzela-Ascoli theorem, a subset of the set of all continuous functions on [0,1]x[0,1] is compact iff it is closed, bounded, and equicontinuous, where equicontinuity of a differentiable function follows from the existence of a global bound on the slope of the function. N is trivially closed. Since all $|h|<1$, the set N is bounded. Finally, because of the bounded derivative, the set is equicontinuous. Therefore, N is compact.

It remains to be shown that T indeed maps N into itself, i.e. that there exists a B such that S_{K+1} will always have derivatives bounded by B as long as S_K is in the set. The bound on the first partial (wrt. v_2) is easy to find:

$\frac{\partial S_{K+1}}{\partial v_2} = \int_0^{v_2} dG(p \mid \ldots) = G(v_2) \leq 1$ because G is a cdf. Therefore, $B \geq 1$.

To show that the partial of S with respect to c_1 is also globally bounded, it is enough to show that the maximum increase in the cdf H(v) of the each of the competitors' v_2 is uniformly bounded. $\frac{\partial}{\partial c_1} \Pr(v_2 \leq v \mid b_{1,K}(v_1, v_2) \leq c_1) < B$

The argument hinges on the notion of an "isobid contour": a line in the (v_1,v_2) space along which the first bid is constant. The constraint imposed by c_1 on other players' valuations is that all their (v_1,v_2) combinations lie below the c_1 isobid contour. Denote the isobid contour corresponding to c_1 as $I(v_2 \mid c_1)$. Two results of implicit differentiation of equation (1.3) ensure that the isobid contour does not move arbitrarily far as a result of a small change in c_1:

1) Slope constraint: $\frac{\partial I(v_2 \mid c_1)}{\partial v_2} = \frac{\partial S(v_2 \mid c_1)}{\partial v_2}$ so from above, $0 < \frac{\partial I(v_2 \mid c_1)}{\partial v_2} < 1$

2) Boundary condition: $\frac{\partial I(0 \mid c_1)}{\partial c_1} = 1$ because when $v_2=0$, S=0.

In particular, because of the boundary condition, the intercept of a isobid cannot shift at a speed greater than unity, and because of the slope constraint, isobids cannot diverge from each other faster than unity. Since the support of v_2 is bounded to [0,1], 1) and 2) together thus imply that $0 < \frac{\partial I(v_2 \mid c_1)}{\partial c_1} < 2$.

Let A_0 be the area under the zero isobid contour $b_1=0$. This area is positive because for low v_1 and high-enough v_2, the bidder bids zero. Please see Figure A1 below for an illustration.

Figure A1: Isobid contours in (v_1,v_2) space, an illustration of the increase in cdf(v) as a result of a small increase in c.

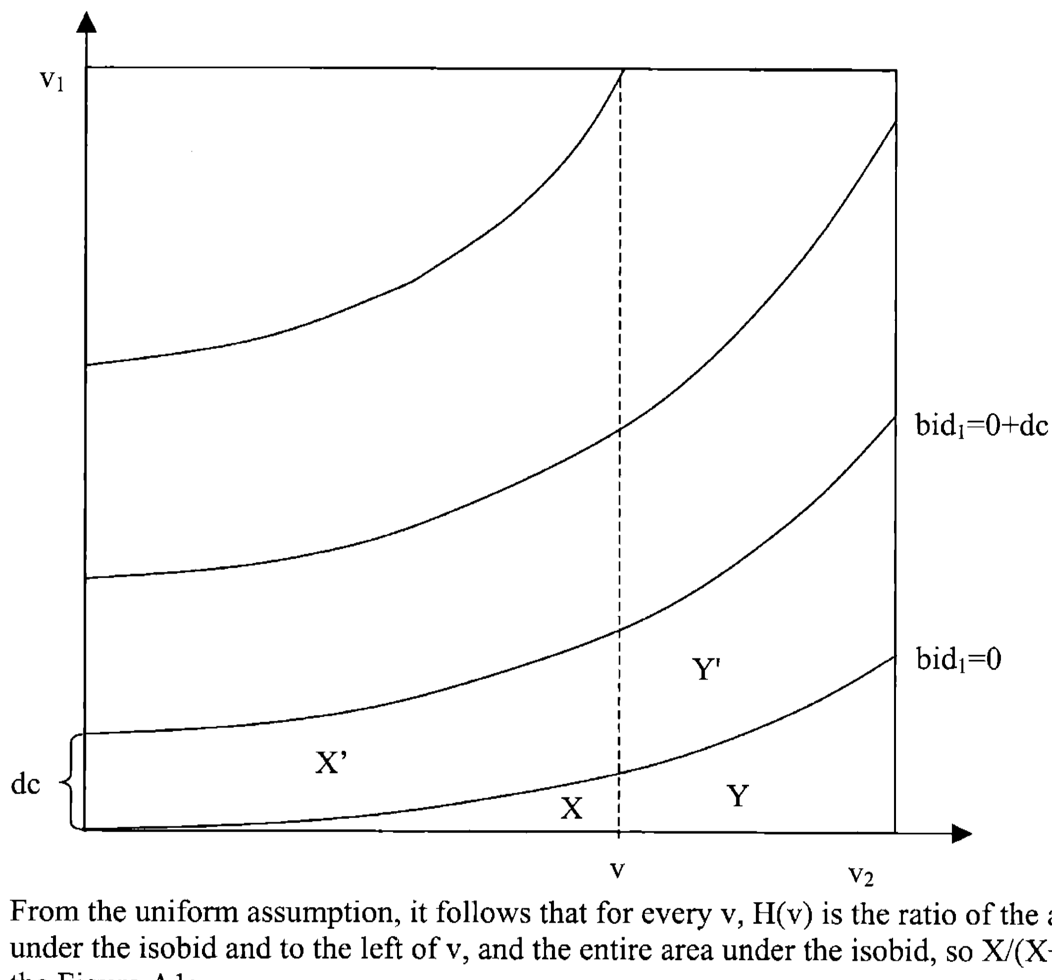

From the uniform assumption, it follows that for every v, H(v) is the ratio of the area under the isobid and to the left of v, and the entire area under the isobid, so X/(X+Y) in the Figure A1:

$$H(v)=\Pr(v_2\le v\mid b_1(v_1,v_2)\le c_1)=\frac{\int_0^{v}\int_0^{I(v_2|c_1)}dv_1dv_2}{\int_0^{v}\int_0^{I(v_2|c_1)}dv_1dv_2+\int_v^{1}\int_0^{I(v_2|c_1)}dv_1dv_2}=\frac{X}{X+Y}$$

Suppose there exists v such that the derivative of H(v) with respect to c_1 has no bound at $c_1=0$. Differentiating H(v) yields: $\frac{dH(v)}{dc}=\frac{X'Y-Y'X}{(X+Y)^2}$ where $0<X'=\int_0^{v}\frac{\partial I(v_2\mid c_1)}{\partial c_1}dv_2<2$ and an analogous constraint holds for Y'. The fact that X' and Y' are bounded leads to a contradiction with the derivative of H(v) having no bound: Since X and Y are fixed positive quantities, the only way the derivative can be arbitrarily large is if X'or Y' are arbitrarily large. This is a contradiction. Therefore, the partial derivative of H(v) is globally bounded, which is sufficient for the partial derivative of S to be globally bounded by some fixed B as well. In case B is actually less than one, let B=1 to obtain a single bound that bounds all derivatives.

From the derivative of isobid being positive and not too large, it is clear that H(v) increases in c_1: the ratio of X'/Y' always exceeds X/Y because an increased c_1 always makes relatively more low v_2's possible.

It thus remains to be shown that T is continuous and differentiable. This follows from the continuity and smoothness of the underlying valuation distribution. *QED*

Proof of Proposition 1.1: The non-decreasing claim follows from a simple differentiation of the equation (1.1) and the fact that cumulative distribution functions are non-negative:

$\frac{\partial S}{\partial v_2}=G_2(v_2)\ge 0$ Given the non-decreasing claim, positivity follows from the boundary condition S(0)=0. Differentiating further, convexity follows from the fact that cumulative distribution functions are nondecreasing: $\frac{\partial^2 S}{\partial v_2^2}=\frac{\partial G_2(v_2)}{\partial v_2}\ge 0$

The function also increases in the highest competitive bid c because as c increases, relatively smaller competitive v_2's are more likely, and the competition hence weakens. *QED*

Proof of Proposition 1.2: S can be bounded from above by some E(x), where E(x) is the expected surplus from facing N-2 bidders drawn from the N-1 lowest bidders of the original N bidders. Let the valuations of those bidders have a distribution function G. Then, E(x) is just: $E(x)=\int_0^{x}(x-p)(N-2)G^{N-3}(p)dG(p)$. The integrand converges to zero exponentially in N as long as G(x)<1. For well-behaved continuous distributions, G(x)<1 whenever F(x)<1. *QED.*

Appendix 1.4: Analysis of Experiment 1 with high outliers included

In this appendix, I present the analysis of the experiment on all the data without excluding the high outliers. All the results are the same in direction, but the standard errors are higher because the high outliers increase the variance of the data.

13 out of the 260 subjects were excluded from the main analysis because they probably did not understand second-price auctions, and put in very high bids (more than $15 on a box of chocolates that costs about $7 in the store, and for which the median bid in this experiment was $4). The $15 cut-off was chosen as a natural break in the histogram of bids, the next highest group of bids was less than $11.

Hypothesis 1: Mean bid in the control condition was $5.46, mean bid in the conditions with future knowledge was $4.72. The difference is $0.74, with one-tail t-test p-value of 0.101. The difference in medians is $1.00.

Hypothesis 2: The slope of the expected surplus function is 1.06 with a standard error of 1.42. Using the linear regression approach implies a slope of $1.43 (error of 1.19), and further controlling for the second bid and familiarity implies a slope of $1.35 (error of 1.10).

Hypothesis 3: The unpredicted main effect of group-size persists, precluding reasonable inference about Hypothesis 3.

Appendix 1.5: Instructions for Experiment 2

Widget Auction Experiment: Instructions

Widgets are new popular gadgets. WiGiTech Inc., the company that makes widgets, found out that everybody would like to have a widget, but that different people value widgets differently. The managers of WiGiTech decided to sell widgets using internet auctions. Before launching the next new line of widgets, WiGiTech is trying to test the new way of selling, and assess its revenue-generating potential.

There will be **20 rounds** of widget sales. In each round, you will be randomly matched with **three other bidders**, and all four of you will participate in **two auctions**. Each auction will sell a different widget, and the first auction will conclude before the second auction starts. Everybody **only wants one widget per round**, so the winner of the first widget will not bid on the second widget in the same round.

We will simulate the fact that different people value widgets differently by secretly **assigning different widget-values** to all participants. Your value of a particular widget is the dollar equivalent of your enjoyment of that particular widget. Whenever you win a widget in an auction, you have to pay the price determined by the auction, so your **payoff from the widget is your value minus the price**. For example, suppose your assigned widget value is $5.45, and you win the widget for a price of $5. Then, you receive a payoff of $0.45. If, however, you win and pay more than your value you **lose money** (Suppose you pay $8 with the above value, then you lose $2.55). When you don't win a widget, you receive nothing. In the end of the experiment, **you will receive the total of all your payoffs** for all the widgets you win.

Everybody's values will be drawn randomly between $4 and $9, and each person will be assigned separate values for the first and second widget within each round. You will learn only your own values for the next two widgets in the beginning of each round, **other people's values will remain secret**. The separate rounds will therefore be completely independent of each other. It is as if the widget market completely restarted in the beginning of every round.

Each auction will be a **second-price sealed-bid auction**. In a second-price sealed-bid auction, all bidders secretly submit their bids for the widget currently on sale. The bids are then secretly compared by the computer. The **highest bidder wins, but pays the second highest bid as the price**. For example, suppose you bid $8 for the widget. If everybody else bids $6, you win and you pay $6 for the widget. However, if one other bidder bids $7.99, you still win, but you now pay $7.99. Highest-bid ties will be resolved randomly, so if one other bidder bids $8 as well, one of you wins at random, and pays $8 for the widget. When at least one other bidder bids more than $8, you lose the widget and get nothing.

In the end of the experiment, **you will be paid the total of your payoffs from all the widgets you won across all rounds**. Note that winning a widget may actually count as a negative payoff if you pay more than your value for it (see example above). Such negative outcomes are added into your total.

Appendix 1.6: Steady-state theory of bargain-hunting on eBay

Proposition 4: The steady-state equilibrium of the K-period look-ahead game is characterized by an expected-surplus function $ES(v_1,\ldots,v_j,\ldots,v_K)$ with the following properties: ES is non-decreasing in all arguments, partial derivatives of ES are all smaller than unity, and they on average decrease monotonically with j.

Proof of Proposition 4: It is clear from the discussion in the theory section that the equilibrium of the K-period look-ahead model is characterized by the expected-surplus function
ES: $\mathbf{R}^K \rightarrow \mathbf{R}$ that satisfies the following Bellman equation:

$$ES(v_0, v_1,\ldots,v_{K-1}) = E_{v_K}\left\{ \int_0^{bid(ES)} (v_0 - p)dG(p) + \delta ES(v_1,\ldots,v_K)[1 - G(bid(ES))] \right\}$$

where

$$bid(ES) = bid(v_0, ES(v_1,\ldots,v_K)) = \max[v_0 - \delta ES(v_1,\ldots,v_K), 0]$$

and where F is the associated unconditional distribution of second-highest bids which follows from whatever the assumptions are on the number of bidders.

To show that the function $ES(x_1,x_2,\ldots,x_K)$ is non-decreasing in all of its arguments, I use a recursive strategy. First, I re-write the Bellman equation in terms of arguments $(x_1,x_2,\ldots,x_K)$:

$$ES(x_1,\ldots,x_K) = E_{x_{K+1}}\left\{ \int_0^{bid(ES)} (x_1 - p)dG(p) + \delta ES(x_2,\ldots,x_{K+1})[1 - G(bid(ES))] \right\}$$

Differentiation with respect to the first argument yields:

$$\frac{\partial ES}{\partial x_1} = E_{x_{K+1}}\{G(\max[x_1 - \delta ES(x_2,\ldots,x_{K+1}), 0])\} \geq 0 \text{ because } G \geq 0.$$

Now, it follows recursively that the function is non-decreasing in the second argument as well:

$$\frac{\partial ES}{\partial x_2} = E_{v_{K+1}}\left\{ \frac{\partial ES}{\partial x_1}(x_2,\ldots,x_{K+1})[1 - G(bid(ES))] \right\} \geq 0 \text{ precisely because } \frac{\partial ES}{\partial x_1} \geq 0$$

And it is clear how this recursion can be carried out through all the other arguments as well, showing that the claim.

Finally, the claim about the relative magnitudes of future partial derivatives follows from:

$$\frac{\partial ES}{\partial x_J}(x_1,\ldots,x_K) = E_{v_{K+1}}\left\{ \frac{\partial ES}{\partial x_{J-1}}(x_2,\ldots,x_{K+1})[1 - G(bid(ES))] \right\} \leq E_{v_{K+1}}\left[\frac{\partial ES}{\partial x_{J-1}}(x_2,\ldots,x_{K+1}) \right]$$

So the claim about relative magnitudes of partial derivatives involves a translation in time and an expectation, and should thus be interpreted as "on average". *QED*

Appendix 1.7: Technical note on imputing time of auction in minutes

For some mysterious reason, eBay only provided the end-date and the duration of each auction. For each bid, we know the date of the bid and the minutes remaining to the end of the auction. In addition, we discern that auction ID-numbers are assigned sequentially. Thus, if auctions A and B both end on the same date and last the same number of days, the auction with the smaller ID ends first.
To put all bids and auctions on a common minute-by-minute time-scale while preserving all the information we have about their order, we do the following for each end-date (called "today" below):

1) For each auction ending today, use all its bids to place a lower and upper bound on its end-minute. The lower bound is the maximum over all bids of:
Minutes_left - 1440*(end_date - bid_date)
Correspondingly, the upper bound is the minimum over all bids of
Minutes_left - 1440*(end_date - bid_date + 1)

2) For each auction-duration (3,5,7, or 10 days), we know the sub-order of auctions by knowing the auction ID. Since, lower bounds should be increasing starting with the first auction and upper bounds should be decreasing starting with the last auction, we "percolate" through the bounds to make them monotonic within each auction-duration.

3) There is no order information to guide the integration across auction-durations. But to generate a consistent ordering, it is enough to sort the auctions by the (percolated) lower bounds. Given than we now get non-monotonic upper bounds, another round of percolation there further tightens the upper bound.

4) The process in 1) – 3) generates a consistently-ordered list of each auction's lower and upper bound on minute end-time. To generate an increasing smooth function of order that fairly closely matches the bounds and maintains the order of the list, we regress the average of the two bounds on a few powers of each auction's order-number and take the normalized predicted value as "pseudo-time". An example of this imputation is shown in the figure.

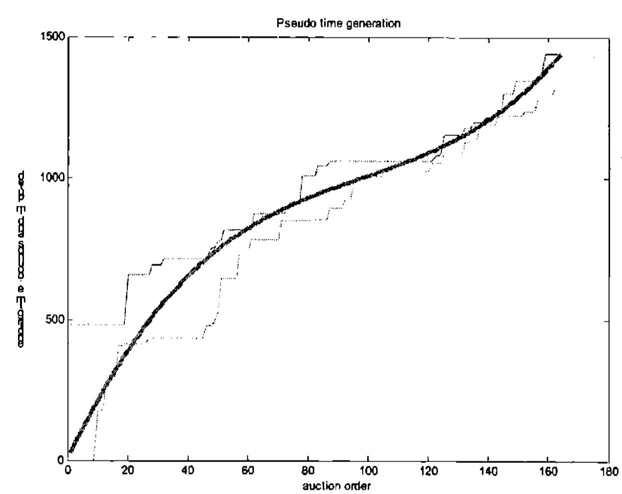

Chapter 2: Dynamics of an auction marketplace with forward-looking bidders: equilibrium analysis

Introduction

When selling multiple units of the same good to several buyers, a traditional seller sets a price, and the buyers take the price as given in making their purchasing decisions. Modern telecommunication technology has enabled an auction-driven marketplace, in which every individual has an effect on every transaction price. The dominant eBay.com site started out in 1995 as a place for exchange of unique collectible items, but has since developed into a marketplace for trading substitute products, especially in categories involving mass-produced consumer products like computers, electronics, music discs, and recently also cars[15]. In such a marketplace, both sellers and buyers can gain from learning and increased strategic sophistication. In a previous paper, I have documented that buyers should and do "bargain-hunt" in anticipation of future options to buy a substitute product by bidding lower than they would in absence of the future options. The multi-unit sellers, on the other hand, have the chance to constantly learn about the momentary profitability of the auction market, and sell strategically to maximize their overall expected profits. In this paper, I analyze a stylized model of the auction-driven market, and characterize the equilibrium that emerges as both sides attempt to secure surpluses from the transactions.

I consider a single monopolist seller, who uses the auction marketplace as the second sales channel besides a stable traditional spot-market. Such a seller faces a constantly changing demand for his products; during any given week, different numbers of bidders with different preferences for the products participate in the auction market. Both official eBay tutorials as well as numerous "how to" books therefore recommend basic market research to sellers. In particular, sellers should gauge the marketability of their items from the sale-prices of the same items in the near past, and constantly adapt to the shifting demand.

[15] The fact that mass-produced consumer goods like digital cameras are sold by auction is an interesting novelty in itself, because before the Internet, this ancient selling format had mostly been used to sell unique goods like artwork, collectibles, and mineral rights.

Since bargain-hunting decreases bids, it may seem that it always hurts sellers and reduces the potential size of the auction-driven market. I will show that sophisticated sellers can take actions to limit the extent of bargain-hunting when the long-term profitability of the auction-driven market is low. Therefore, bargain-hunting is a self-regulating phenomenon that turns itself off just when the very opportunity to trade is threatened. However, this does not mean that bargain-hunting persists in equilibrium only when the seller always sells. It is possible for bargain-hunting to co-exist with a discriminating seller, because there are informational benefits associated with bargain-hunting buyers. By bargain-hunting, buyers reveal more to the seller about their valuation and number, so the seller can in turn make better selling decisions. This future benefit sometimes outweighs the current cost of lower prices, and the seller may thus even want to sometimes encourage bargain-hunting.

To maintain tractability, I use a very simple stylized model. Before describing the modeling setup in detail, let me outline its main features. I assume that the outside spot market is stable, and I normalize the seller's spot-market profits to zero. Given this assumption, the model is only interesting if the auction-market prices can be both positive and negative. I thus assume that the bidders in each period could be of either the negative "Low" type or the positive "High" type. As a model for each auction, I use the second-price sealed-bid auction because in such an auction, bidder's strategies are simple functions of bidder values, and dominant strategies often exist, helping limit the extent of convoluted strategic considerations. To limit the complexity of the optimal bargain-hunting strategy, I limit the bidder life-span to two periods, and I assume the demand-side is being refilled with a steady number of bidders every period. This assumption makes sense if bidders use the auction marketplace to fill time-sensitive needs. For example, if I want to buy an electronic gadget soon, I may try eBay for a week, but if I remain unsuccessful, I go and buy it in the store. The assumption about bidder life-spans also delineates what can the seller learn about demand: the seller is facing both new and old bidders in each period, and he can use past prices to deduce the number and type of the old bidders. Based on his learning about the old bidders, the seller either sells or not. Based on the seller's strategy, the buyers either bargain-hunt or not. How do all these possibilities work together?

Literature review

Models of optimal selling in sequential auctions are relatively rare in the literature, especially when the buyers are assumed to be strategic and forward-looking. All models of bidding in sequential auctions assume that the seller exogenously sells in every future period (for example Milgrom & Weber 1982b, Engelbrecht-Wiggans 1994, Gale & Hausch 1994), and conclude that the bidders therefore shade their bids down. This paper contributes to the literature on sequential auctions by analyzing the best response of the seller to such strategic bidding, and characterizing an equilibrium of the buyer-seller game.

In the proposed model, the seller decides each period whether to sell in the auction market or in some other stable spot-market. Besides the selling decision itself, the seller's strategy in the auction market can also involve other instruments that are affected by strategic bidding, most notably the reserve price. McAfee and Vincent (1997) consider the impact of long-lived strategic buyers on optimal reserve prices a single object is being sold. They find that when the seller cannot commit to never re-selling the unsold object, the effectiveness of reserve-prices is greatly diminished because buyers anticipate a resale in the case that their bids do not meet the reserve price. In response, the "sequentially optimal" reserve price is smaller and about ninety percent less effective than the classic optimal reserve price in a single auction analyzed by Myerson (1981). Therefore, at least in the case or reserve prices, the seller should act differently when facing forward-looking strategic bidders who live longer than a period.

If it is costly for the bidders to enter the auction, an even further reduction in the overall benefit from reserve prices can be predicted from the results of Engelbrecht-Wiggans (1993). He finds that if a reserve price turns even one bidder away, the extra profit it provides is not worth the profit foregone by having fewer bidders. Taken together, the results of Engelbrecht-Wiggans and McAfee and Vincent suggest that in realistic auction-driven markets, reserve prices are not nearly as beneficial to the seller as it may seem from the simple case of a single isolated auction and a captive audience. I therefore abstract away from reserve prices and focus on the selling decision while modeling each auction as an auction without a reserve price.

While papers on optimal selling in sequential auctions are rare, a lot of work has been done on optimal sequential posted-price selling to price-taking buyers. The most related paper is the model of Besanko and Winston (1990) of a monopolist seller facing strategic consumers. Besanko and Winston find that the fact that consumers are strategic and forward-looking has a dramatic effect on the optimal sequence of prices when the seller cannot commit to future prices. In particular, the seller cannot use "price-skimming", i.e. high initial prices, as effectively as if the buyers were myopic. Fundamentally, this is because analogously to forward-looking bidders who reduce their bids in sequential auctions, the forward-looking buyers refuse to buy at a given price because they rationally expect the future price to be lower. In other words, the demand of forward-looking buyers is more elastic than that of myopic buyers. Analogously to the situation in McAfee and Winston (1997), the sellers ability to intertemporally price-discriminate is reduced when the bidders become strategically forward-looking, a principle that goes back to the seminal ideas of Coase (1972).

One way to contrast the model in this paper with the models discussed so far is that in the proposed model, the sellers can effectively respond to the threat of forward-looking strategic bidding. As in McAfee & Vincent or Besanko & Winston, strategic bidders limit the ability of the auction seller to extract profits. On the other hand, this paper also shows that sellers with outside options can limit the ability of strategic bidders to extract surplus. In particular, strategic withholding of future supply diminishes the power of bargain-hunting, and sometimes even makes it not profitable at all. This consideration makes the resulting equilibrium more interesting.

To obtain non-trivial effects of strategic bidders on the optimal seller strategy in sequential auctions, it is not necessary for the bidders to be long-lived. Vulcano, van Ryzin and Maglaras (2002) (VRM) study a fixed-capacity seller who sells multiple units of a good overtime to unit-demand strategic buyers who only live for one period. In each period, the seller collects bids from all buyers, and decides how many of the remaining units of the good to sell that period. VRM discover that it is not optimal for the seller to first find the optimal number of units to sell k^* given the opportunity cost of remaining inventory, and then just sell to the k^* highest bidders for the price equal to the (k^*+1)-th highest bid. Given such a seller strategy, it is no longer dominant strategy for the bidders to bid their valuation as it would be in a simple (k+1)-th price Vickrey auction. Instead, there is an incentive to overbid and make the seller sell more units than would be optimal for him. Therefore, VRM provide an example of strategic selling in a sequence of auctions having an effect on the optimal bidding strategy. They also find that there exists a simple modification to the (k^*+1)-th auction that actually restores the truth-telling property of the Vickrey-like mechanism, and that there is thus an equilibrium of the buyer-seller game, in which the buyers bid their valuation. It is not clear, however, whether such an equilibrium would be possible if the buyers lived for more than one period.

This paper is also related to the channel literature because the seller is constantly choosing between a fixed outside option and the auction-market. Wang (1990) has studied the choice between posted-price and auction selling of a single object, and concluded that ceteris paribus, auctions become preferred as the dispersion of the bidder valuations increases. The distribution assumed here is just a discrete {Low,High} distribution, but Wang's result seems to hold because as difference High-Low increases, the seller is more and more likely to sell in the auction market versus the fixed outside market. Because the outside option assumed here is so simple, any more detailed conclusions are likely to be of limited value.

Model

- Time is discrete and discounted at a factor δ per period[16] by both sellers and bidders.
- Risk-neutral bidders live for two periods. Exactly two bidders start each generation, and the winner drops out of the game (unit-demand)[17]. In the end of their life, all remaining bidders drop out of the game. Each generation of the bidders is of valuation type **Low** or **High**, where valuation is the utility of the product to a single bidder measured in dollars. Probability of Low is **p**. Bidders' valuations are thus perfectly correlated within a generation, but independent across generations. The within-generation correlation increases the demand fluctuations and simultaneously helps the seller to learn about the type of remaining bidders from past prices. Learning would be possible but harder with within-generation independence, and the main qualitative results of this paper would go through. The case of uncorrelated bidders is discussed in detail later in this paper.
- A monopolist seller is endowed with an object to sell each period. In each period, the seller decides whether to sell the item on eBay or on affixed-price spot market. Normalizing the spot-market price to zero, non-trivial behavior occurs only when the relative profitability of selling on eBay vis-à-vis the spot-market is ex-ante uncertain, namely if H>0>L. The eBay auction is modeled as a second-price auction without a reserve, bid-ties get resolved by randomization. In each period, the seller thus faces two "new" bidders and one or two "old bidders". When the seller decides not to sell on eBay, he is sure to face two old bidders a period later, but the only information available about those bidders' valuation is the prior: p·Low+(1-p)·High. I assume that the seller starts in exactly that situation in period one, and then uses the Bayes theorem in all subsequent periods to update his beliefs about the number and type of old bidders. Starting with a different correct belief would simply involve an initial ramp-up before the game analyzed here.

---

[16] A time-period relevant for internet auctions like eBay is a week.

[17] Two bidders per generation is the maximum non-trivial case given the assumption of perfect correlation. With three or more bidders, there is always a tie at the high bid, so there is never any buyer surplus, and therefore no bargain-hunting.

Bidder strategy

Most of the time, the bidders bid just like they would in a single isolated auction. In their second period, the surviving bidders have a dominant strategy to bid their valuation. In their first period, Low-valuation bidders bid valuation because the least competitive bid by old bidders is L, and bidding anything less than L would have no chance of winning.

The only time bidders bid anything else than their valuation is when they expect to make a positive surplus in the future should they lose the current auction. Then, High-valuation new bidders bargain-hunt:

$$\text{Bid(H,now)=H-E[surplus(next period)|lose now]} = H - a < H \tag{2.1}$$

Two conditions need to hold for the expected future surplus of a High new bidder to be positive given that he loses the current auction:

1. a nonzero chance that the seller will offer a product for sale in the auction market in the next period.
2. a nonzero chance that the other bidder of the same generation is not present in the next period, i.e. that he wins the current auction

Therefore, the new-bidder beliefs concerning the type of current old bidders play a key role in making bargain-hunting possible. I assume that they hold prior beliefs about the valuation of their current competition from the old bidders. Then, the second condition for the expected future surplus is satisfied. I will investigate other possible assumptions after analyzing the whole buyer-seller model.

Given the nonzero chance of a new High bidder winning in the current period, the strategy of the new High bidders depends only on the chance σ that the seller sells in the next period. If $\sigma>0$, the expected surplus from the next period is positive because the High-valuation bidder might be the only remaining one in his generation, and the new generation may be Low. When, on the other hand, $\sigma=0$, there is no surplus to be had in the next period, and Bid(H,now)=H. Note that the σ probability is not a model parameter, it is determined by seller behavior.

The actual bid-decrement arises from the game played between two new High bidders. I focus on symmetric equilibria, in which bid is less than valuation. Suppose the

High-type new bidders bid H-a. The only way they can make a positive surplus is if there are no High old bidders, because when there is at least one old High-type bidder, the new bidders who bid H-a get zero total surplus since they both advance to the next round and bid H there. Therefore, both bargain-hunting bidders assume (and hope) that the old competition is of type L. Then, one of the new bidders wins the object by bidding H-a, and drops out of the game with a surplus of a. The remaining bidder then wins the next auction and collects the expected surplus p(H-L)+(1-p)a one period later if the seller sells in that period. Bidding H-a in the first period is then a symmetric pure-strategy equilibrium if surplus today equals expected surplus tomorrow, preventing profitable deviations:

$a = \delta\sigma[\, p(H-L)+(1-p)a]$, so $a[1-\delta\sigma(1-p)] = \delta\sigma p(H-L)$, and therefore $a < (H-L)$

Therefore, $L < H-a < H$, so bargain-hunting leads to prices between L and H. When bidders have more information about the competition and about seller's strategy, bargain-hunting may change both quantitatively and qualitatively. I will investigate this possibility after analyzing the seller model.

Suppose new High bidders always bargain-hunt. Here is a summary of prices arising from the bidder model:

Table 2.1: Prices in the auction market		New bidder type	
		L (p)	**H** (1-p)
One old bidder	**L** (p_{1L})	L	H-a
	H (p_{1H})	L	H-a
Two old bidders	**L** (p_{2L})	L	H-a
	H (p_{2H})	H	H

Seller strategy

The seller learns about the type and number of old bidders from prices. There are four possible combinations of old-bidder number and type: $\{1L,1H,2L,2H\}$, so seller's beliefs are characterized by a vector of probabilities (p_{1L}, p_{1H}, p_{2L}, p_{2H}). If the seller decides to sell on eBay, he updates his beliefs using the observed price using the Bayes rule:

Price =

L: then the new bidders must have been Low, so:

$$(p_{1L}, p_{1H}, p_{2L}, p_{2H}) \to \left(\frac{\frac{2}{3}p_{1L} + \frac{1}{2}p_{2L}}{p_{1L} + p_{1H} + p_{2L}}, 0, \frac{\frac{1}{3}p_{1L} + \frac{1}{2}p_{2L} + p_{1H}}{p_{1L} + p_{1H} + p_{2L}}, 0 \right)$$

H-a: then the new bidders must have been High, so:

$$(p_{1L}, p_{1H}, p_{2L}, p_{2H}) \to \left(0, \frac{p_{1L} + p_{2L}}{p_{1L} + p_{2L} + p_{1H}}, 0, \frac{p_{1H}}{p_{1L} + p_{2L} + p_{1H}} \right)$$

H: then there must have been two old bidders, and no information beyond the prior is available about the type of new bidders. It is certain, however, that neither of the new bidders dropped out:

$$(p_{1L}, p_{1H}, p_{2L}, p_{2H}) \to (0, 0, p, 1 - p)$$

By assumption, the seller starts with the correct belief (0,0,p,1-p), and he returns to that belief whenever he does not sell in a period. From above, the seller also returns to this starting belief whenever he observes the price H. Before studying the seller model further, I present three observations about the nature of bidder learning from prices.
The first observation about the seller's learning process is that learning about bidder type is more accurate when prices are lower. Starting with (0,0,p,1-p), learning occurs when Price=H-a or Price=L. On the other hand price=H contains no information about the type of new bidders. In a continuous model, the same should intuitively hold about learning from prices. The lower the price, the lower the variance of the posterior distribution of

the valuations of the surviving bidders. It would be an interesting extension to investigate the variance of the payoff-relevant posterior distribution, namely the distribution of the maximum of the valuations of the surviving bidders.

The second observation about the seller's learning process is that learning about new bidders is a function of the knowledge of the current bidders. Therefore, the seller's knowledge evolves over time, and it is not a zero-order process.

Finally, the third observation about the seller's learning process is that it is particularly enabled by the phenomenon of bidder bargain-hunting. The seller actively separates the implied meaning of the High price on one hand and the intermediate H-a price on the other. In other words, if a=0, the seller's learning from prices would be different and less accurate.

The seller's belief-vector is a knowledge state in the sense that it summarizes learning to date, and the seller's learning describes the transitions among states over time when the seller continues selling n the auction market. When the seller sells in the spot market for a period, he transitions back into the initial state. Therefore, for any strategy $(p_{1L}, p_{1H}, p_{2L}, p_{2H}) \rightarrow \{\text{sell, not sell}\}$, the seller moves through the state-space of all possible belief vectors. While in general it would be necessary to characterize the optimal behavior in a continuum of states, the seller's information in this particular game can only be in of four profit-relevant states:

1) **Q**: (0,0,p,1-p) , the uncertain starting state
2) **1H**: (0,1,0,0) , the state when it is certain there is one High old bidder
3) **2H**: (0,0,0,1) , the state when it is certain there are two High old bidders
4) **L**: (q,0,1-q,0) , the state when it is certain there are no High old bidders

The transitions among these states are illustrated by the flow-diagram in Appendix 2.1. Note that all states in which Prob(Old=H)=0, namely all the states of the form (q,0,1-q,0) are equivalent with respect to seller's current profit (seller gets pL+(1-p)(H-a) if he sells) and the seller learning (price L means (q,0,1-q,0), price H-a means (0,1,0,0)) . Therefore, I can pool all these states into a single "L" state in the analysis of the seller's profit-maximizing strategy.

A seller's strategy therefore amounts to selling in the auction marketplace as a function of the state. The seller maximizes net present value of profits, so the optimal steady-state profit function $\Pi: \{L,Q,1H,2H\} \rightarrow R$ is characterized by the following set of Bellman equations in general:

General Bellman Equations With Bargain-Hunting

$\Pi_L = \max\{\delta\Pi_Q \text{ , } pL+(1-p)(H-a) + \delta p\Pi_L + \delta(1-p)\Pi_{1H}\}$

$\Pi_Q = \max\{\delta\Pi_Q \text{ , } p^2L+(1-p^2)H-p(1-p)a + \delta p^2\Pi_L + \delta(1-p)\Pi_Q + \delta p(1-p)\Pi_{1H}\}$

$\Pi_{1H} = \max\{\delta\Pi_Q \text{ , } pL+(1-p)(H-a) + \delta p\Pi_L + \delta(1-p)\Pi_{2H}\}$

$\Pi_{2H} = \max\{\delta\Pi_Q \text{ , } H + \delta\Pi_Q\} = H + \delta\Pi_Q$

Moreover, all profits must be greater than zero, the spot-market price.

From the Bellman equation for Π_{2H}, it is clear that the seller always sells in the 2H state. It is also clear that the seller eventually reaches the uncertain state Q, so strategies in which the seller does not sell in Q essentially mean exiting the auction market completely. Therefore, I focus on the region of the parameter space, in which the seller sells in Q.

Besides the trivial strategy of always selling in the auction market when all states are associated with positive profit, the seller can sell in two different ways. Since intuitively, $\Pi_{2H}>\Pi_{1H}$, selling in the L state implies selling in the 1H state. Therefore, the seller either strategically withholds supply from the auction marketplace in the L state only, or "pulses" by withholding supply in both the L and the 1H states, effectively selling every other period. The key aspect of the seller strategy is therefore whether or not the seller sells in the state 1H.

Equilibrium in the game between the seller and the bidders

While pulsing is clearly incompatible with bargain-hunting, the strategy of strategic withholding in the L states supports bargain-hunting in equilibrium because the seller does sell in the only buyer-profitable state 1H. The main theorem of this paper characterizes the two different equilibrium strategies in terms of the overall model's parameter space. To characterize the seller's optimal policy, it is useful to introduce notation for the short-run expected revenues of the seller in the Q state and in the L/1H states when bidders are bargain-hunting, ER_Q ,and ER_L respectively:

$$ER_Q = p^2L+(1-p^2)H-ap(1-p) \; > \; ER_L = pL + (1-p)(H-a)$$

The following theorem then characterizes the equilibrium of the game between the seller and the bidders.

Theorem 2: (please see Appendix 2.2 for proof)

The equilibrium of the auction market depends on the profitability of the auction-market as follows:

- When the auction-market is relatively more profitable, but not enough for the seller to abandon the spot market completely, namely when $-\delta p(1-p)H < ER_L < \frac{\delta(1-p)H}{1+\delta-\delta p}$, the seller learns from prices to withhold supply from the auction market whenever the old bidders are of the Low type, and the bidders bargain-hunt.
- When the auction-market is relatively less profitable such that $ER_L<-\delta p(1-p)H$, the bidders do not bargain-hunt. The seller either uses the pulsing strategy, or also sells in other informational states associated with the new learning environment without bargain-hunting.

The profits available to the seller when he sells on the auction market:

Pulsing: $\Pi_Q = \frac{ER_Q + ap(1-p)}{1-\delta^2}$

Strategic Withholding in L state: $\Pi_Q = \frac{ER_Q + \delta p(1-p)[ER_L + \delta p(1-p)H]}{(1-\delta)(1+\delta p+\delta^2 p(1-p))}$

For completeness, it should be noted that when $ER_Q + ap(1-p)<0$, there are no sales at the auction market because even the pulsing strategy delivers a negative profit, and when

$\frac{\delta(1-p)H}{1+\delta-\delta p} < ER_L$, the seller sells only on the auction market while the bidders, of course, bargain-hunt.

Theorem 1 is illustrated in the graph below. While it has been beneficial for exposition to carry around both and L and H in general, it is WLOG to set L to −1. The only remaining model-parameters are H, p, and δ. Fixing one of these parameters, it is possible to illustrate the regions of the parameter space corresponding to different equilibrium situations. I fix δ at 0.95, and focus on the (H,p) space.

Figure 2.1: Auction-market equilibria in (H,p) space

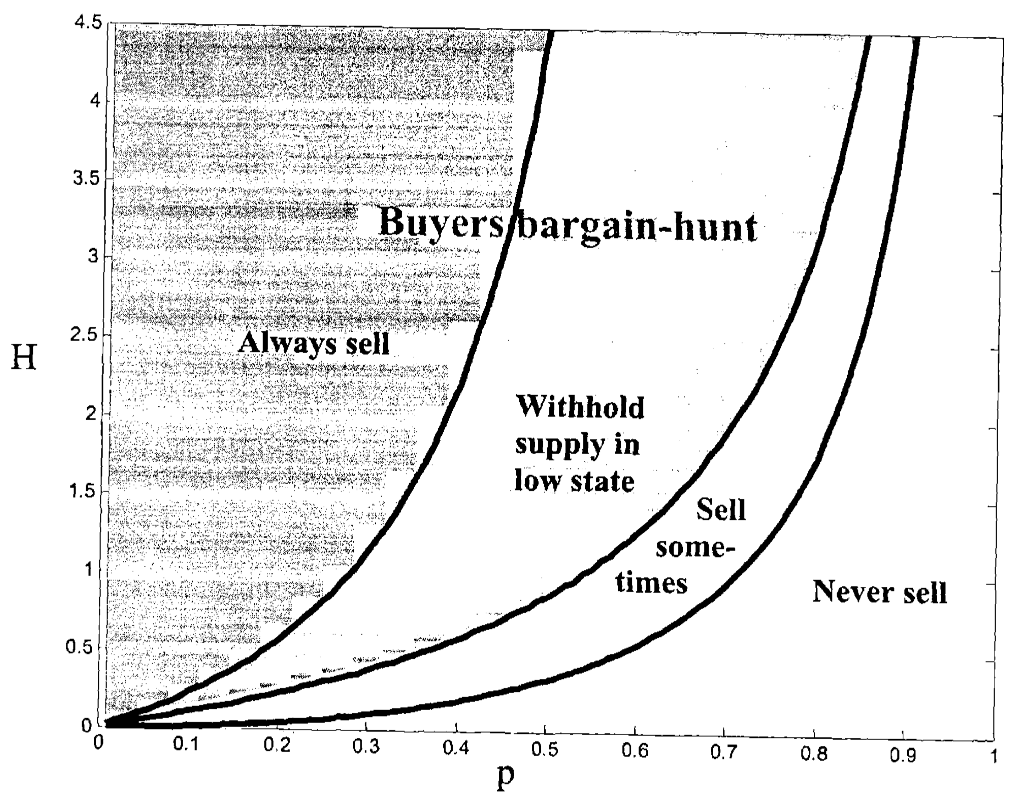

In Figure 2.1, the buyers bargain-hunt in the gray region while the different seller strategies are delineated by black boundaries. The three boundary curves correspond to the inequalities in the Theorem.

Intuition for Theorem 2

In a model where learning is sometimes available as a side-effect of selling, the question is not so much when to sell, but rather when to stop selling. As long as the seller always sells, the bidders obviously bargain-hunt. The seller's ability to learn from prices plays a critical role in making bargain-hunting possible in equilibrium even when the seller does not always sell. The bidders continue bargain-hunting as long as the seller sells in his 1H knowledge-state, namely the state in which the old bidder makes a surplus. Why would the seller want to sell in that state and potentially give the (H-L) surplus to the old bidder? The reason is that while short-term payoffs may be relatively low, the future payoffs can outweigh the short-term payoffs precisely because learning from prices enables the seller to sell in the following period only if he finds himself in state 2H. In fact, the $ER_L > -\delta p(1-p)H$ condition is a direct evaluation of the net tradeoff the bidder is facing at this point: by losing ER_L now, he can make $\delta p(1-p)H$ in expectation before returning to the Q state, accounting for the fact that he reaches the Q state a period later than he would should he withhold supply and not sell in the 1H state.

When the seller does not find it profitable to sell in the 1H state, the buyers in turn do not find it profitable to bargain-hunt. The reason is that with the seller not selling in L or 1H, there are no two successive periods that both allow the bidders to make a surplus. In particular, the only two successive periods in which there is a sale are both in state Q, but the transition from Q to Q ensures that old bidders earn no surplus in the second period. Therefore, when selling in 1H is not profitable to the seller, he has a credible threat that ensures bidders do not bargain-hunt.

It is not immediately clear what happens in a world without bargain-hunting, because the seller's ability to learn from prices decreases, and the number of payoff-relevant states becomes infinite. Please see Appendix 2.3 for an illustration of the situation. It is clear that as the profitability of the auction market decreases further, the seller eventually exits completely. The least enthusiastic seller, who still participates in the auction market, uses a pulsing technique, always waiting one period for the demand side to re-fill with two old bidders. In some sense, this suggests that knowing the number of old bidders is more important for good selling decisions than knowing their type. This is because a single old bidder, no matter how high his valuation, does not influence the price when the new bidders are Low. Going to two old bidders thus strictly increases the profit more often than increasing a single old bidder's valuation.

Effects of bidder learning

The realistic assumption about the new bidders being unaware of the seller's state does not play a key role in the main result. Suppose the bidders can interpret prices as well as the sellers can (despite having shorter life-spans), and thus know which state the seller is in. Would such bidders still bargain-hunt? Not as often. In particular, the bidders would bargain-hunt only in states that precede 1H, namely Q and L. However, there would be no bargain-hunting in 1H itself, because the expected future surplus of bargain-hunting bidders who are new when the seller is in 1H is zero. Therefore, bargain-hunting in Q and L would satisfy $a' = \delta p(H-L) < a$. The profits and cutoffs would have to be re-calculated, but the seller strategy would be qualitatively the same as in Theorem 1. Overall, this situation would help the seller because his ability to learn from prices is not degraded, while there is less bargain-hunting. Therefore, we might expect that the seller would like to announce his state. This intuition might generalize as follows: when an eBay seller is uncertain about demand, he gains from announcing that he has future units for sale. This promotes bargain-hunting and hence learning. When, however, the seller is not uncertain, he waits until after the first sale to announce another potential sale. This way, the seller encourages bargain-hunting only when its benefits outweigh its costs.

What about reserve prices? Do they erase all these results?

Yes, reserve prices are useful insurance against the low-profit outcomes the seller is trying to avoid. However, research by McAfee and Vincent suggests that reserve prices are of limited effectiveness in sequential auctions (please see the Literature Review for a discussion). Even if reserve prices were incorporated, this work is still relevant because there are still benefits to withholding supply for a period, and there still is bargain-hunting. Simply setting reserve price equal to the spot-market price just cuts off the bottom of the profit-distribution, but leaves the top still there and interesting. And it is the top that influences profits.

What about multiple sellers? Doesn't free learning from other sellers' sales erase the tradeoff between selling and learning? As long as the information from own sales is better than information from others' sales, the tradeoff operates. Whenever the seller sells special goods, this should be true. Also, seller can better control his own selling, and hence can interpret the determinants of prices better. Extending to multiple sellers would be interesting, there would probably be mixing in equilibrium of the sellers' game. It is clear that selling has a positive externality in the information the final price provides about the state of demand. One could hypothesize that there is therefore free-riding and not enough selling in auction markets.

What if the two bidders within each generation are not perfectly correlated?

The assumption of perfect within-generation correlation contributes greatly to the simplicity of the model, but it is not necessary for the qualitative conclusions. To demonstrate that the model extends even to situations when the bidders are each drawn at random from some distribution, I will now assume that the bidders in each generation are actually independent.

Suppose each bidder can be High or Low, the probability of Low is p, and two new bidders are drawn independently in each period. Then, the price-table has one more column H/L to capture the situation when the new bidders are of different types, and one more row to make the same possible for the Old bidders. As in the original model, the seller state 1L is still payoff-equivalent to all convex combinations of states 1L and 2L because when the new bidders are H/L and the old bidders $q \cdot 1L + (1-q) \cdot 2L$, the high new bidder wins with certainty, and all remaining old bidders are thus Low. Moreover, the seller states 1H and H/L are also payoff-equivalent, because the additional Low old bidder in H/L does not affect learning or payoffs. Call this class of states 1H. Therefore, the price-table over payoff-relevant states is as shown in Table 2.2

Therefore, the difference in the price-table from the seller's perspective is merely the addition of the H/L column. Not surprisingly, this degrades the information that can be obtained from prices, but it does not degrade it completely. A Low price yesterday still means that today's bidders are all of the Low type because when the new bidders yesterday were H/L and the price was L, the High new bidder won and exited the market. A High price still sends the seller to the uncertain Q state regardless of where it occurs.

Table 2.2: Prices in the auction market with independent bidders		*New bidder type*		
		L (p^2)	**H/L** $2p(1-p)$	**H** $(1-p)^2$
Payoff-relevant seller states	**L** (p_L)	L	L	H-a
	1H (p_{1H})	L	H-a	H-a
	2H (p_{2H})	H	H	H

The only difference in learning from prices is that when the seller in the Q state observes the intermediate H-a price, he cannot immediately infer the state 1H, but only a mixture of the 1H and 2H states. The fundamental reason for this complication is that the unknown old bidders in the Q state can actually be H/L, in which case both new High bidders advance to the next round. The exact Bayesian updating is:

Price = **H-a**: then the new bidders must have been High or H/L, and so the next period's old bidders must belong to the 1H state unless the old bidders were actually in the 1H state (really H/L), and the new bidders were actually both High.

$$(p_L, p_{1H}, p_{2H}) \rightarrow \left(0, \frac{(1-p)^2 p_L + 2p(1-p)p_{1H}}{(1-p)^2(p_L + p_{1H}) + 2p(1-p)p_{1H}}, \frac{(1-p)^2 p_{1H}}{(1-p)^2(p_L + p_{1H}) + 2p(1-p)p_{1H}}\right)$$

Therefore, from the Q state, the H-a price sends the seller to a mixed state M0 between 1H and 2H, fully characterized by the probability of 2H p_{2H}^0:

$$p_{2H}^{0} = \frac{(1-p)^2 2p(1-p)}{(1-p)^2\left(p^2 + 2p(1-p)\right) + 2p(1-p)2p(1-p)} = \frac{2-2p}{2+3p}$$

The most important aspect of this state is not the exact probability of 2H, but the fact that the probability of Low state is zero. If there is no chance the seller is in the L row of the price table, a price of H-a leads to a mixed state M1 with a fixed probability $p_{2H}^{1} = \frac{1-p}{1+p} > p_{2H}^{0}$. In other words, today's p_{1H} drops out of the updating equation. Therefore, the seller can only be in five payoff-relevant states in this game: L,Q, 1H, M0 and M1, one of which, the 1H state, is only accessible when the seller always sells.

Consider the more interesting situation, in which the seller does not sell in state L. Since the continuation from the M0 and M1 states is qualitatively the same, the fact that $p_{2H}^{1} > p_{2H}^{0}$ ensures that the M1 state is more lucrative than the M0 state, and selling in M0 thus implies selling in M1. Moreover, M1 is only accessible through M0, so the seller either sells in both M0 and M1 or uses the pulsing strategy. Whenever withholding supply in M0 is a credible threat, the bidders will not bargain-hunt, and the main qualitative conclusion of Theorem 1 will hold.

The Bellman equations for the case, in which the sellers sells in M0 but not in L are:

$\Pi_Q = p^3(4-3p)(L+\delta^2\Pi_Q)+p(1-p)^2(2+3p)(H-a+\delta\Pi_{M0}) +(1-p)^2(H+\delta\Pi_Q)$

$(2+3p)\Pi_{M0} = 5p^3(L+\delta^2\Pi_Q)+5p(1-p^2)(H-a+\delta\Pi_{M1})+(2-2p)(H+\delta\Pi_Q)$

$(1+p)\Pi_{M1} = 2p^3(L+\delta^2\Pi_Q)+2p(1-p^2)(H-a+\delta\Pi_{M1})+(1-p)(H+\delta\Pi_Q)$

$\Pi_{M0}>\delta\Pi_Q>0$

$\Pi_L=\delta\Pi_Q>p(2-p)(L+\delta\Pi_L)+(1-p)^2(H-a+\delta\Pi_{1H})$ where Π_{1H} is the state with one old H bidder.

The solution of $\{\Pi_Q, \Pi_{M0},$ and $\Pi_{M1}\}$ is extremely algebraically involved, but an informal argument can be for the claim that withholding supply in M0 eventually becomes a credible threat as the profitability of the auction-market shrinks.

Discussion

Bargain-hunting is a self-regulating phenomenon. When a stream of buyers goes through a sequential auction-driven marketplace controlled by a single seller, the buyers do not always bargain-hunt. The seller always has the option to sell elsewhere, and the relative profitability of the auction determines whether the market is in a bargain-hunting equilibrium or not. To decide when to sell where, the seller observes recent prices and learns about both the number and type of the remaining buyers. When the auction marketplace is sufficiently relatively profitable, the seller sells there often enough to make bargain-hunting worthwhile. When, on the other hand, the auction marketplace is not too profitable to the seller, bargain-hunting becomes unprofitable to the buyers as well by virtue of a change in the seller's strategy. For example, in the extreme case, the seller uses a pulsing strategy, spacing sales in the auction marketplace enough from each other to let the demand "refill" completely to the maximum possible number of bidders. In such a world, short-lived buyers have nothing to gain from bargain-hunting. Therefore, bargain-hunting switches itself off just when the existence of the marketplace is threatened by low seller profits, and it therefore does not have an impact on market size.

This self-regulation is made possible by the extra information contained in prices whenever bidders bargain-hunt. Bargain-hunting has not only the obvious cost to the seller, but also a benefit in that it improves the seller's ability to learn from prices: the prices are more dispersed, and the underlying buyer types thus more easily discernible. Because of this benefit, bargain-hunting persists in equilibrium even when the seller could perfectly prevent it by threatening to strategically withhold future supply. Such a threat only becomes credible when the learning benefit is not large enough, which in turn happens when the profitability of the auction marketplace decreases further. It is exactly this credible threat that helps the seller turn bargain-hunting off when the very existence of the auction-market is threatened.

It is informative to consider the properties of the basic building-block of the model - seller learning from past prices. When sellers try to learn about the new bidders who just arrived to the auction marketplace and will be around in the future, knowing who the old bidders are helps the seller deduce information about the new bidders.

Therefore, the learning is not a zero-order process, and seller learning is well captured by an evolving knowledge-state. Without bargain-hunting, the seller can be in infinitely many knowledge states, and the optimal selling strategy is thus hard to classify. When, on the other hand, the buyers bargain-hunt, they reveal more about their types, enough to reduce the relevant space to just four knowledge-states. It is not clear how this would generalize to continuous type-spaces, but bargain hunting makes the seller's strategy-optimization substantially easier in the present model.

Appendices of Chapter 2

Appendix 2.1: Flow-diagram of model with bargain-hunting

Always sell:

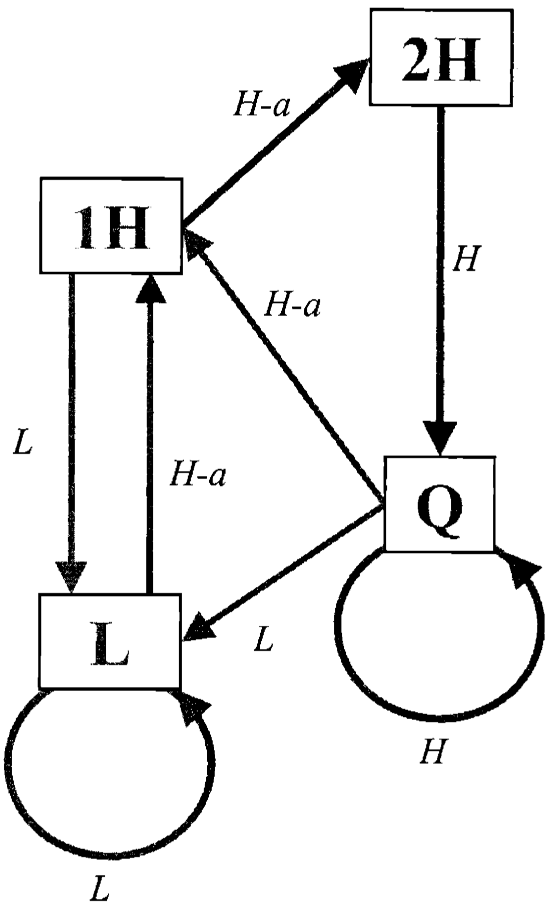

Not sell in L state:

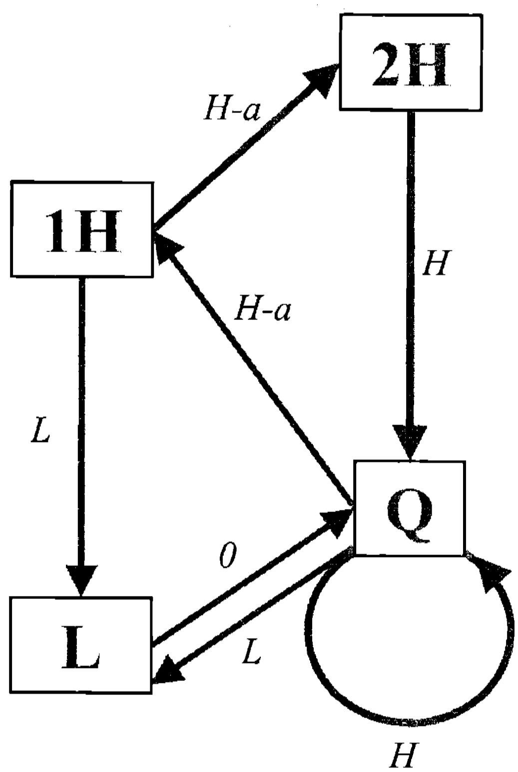

Appendix 2.2: Proof of Theorem 2

First, observe that bargain-hunting persists iff the seller sells in the 1H state. Suppose otherwise, i.e. bidders bargain-hunting and the seller not selling in 1H. Then, the seller does not sell in state L and never reaches the state 2H, so the only two sales right after each other are in two successive Q states, but the only way that can happen is when the two bidders in the first Q state are high, and that situation leaves no surplus to the new High bidders.

Therefore, suppose first that there is selling in the 1H state, and so there is bargain-hunting. Then, the equations reduce to:

$$\Pi_Q = p^2L+(1-p^2)H-p(1-p)a + \delta p^2\Pi_L + \delta(1-p)\,\Pi_Q+\delta p(1-p)\,\Pi_{1H}$$
$$\Pi_L = \max\{\delta\Pi_Q\,,\, pL+(1-p)(H-a) + \delta p\Pi_L + \delta(1-p)\Pi_{1H}\}$$
$$\Pi_{1H} = pL+(1-p)(H-a) + \delta p\Pi_L + \delta(1-p)(\,H + \delta\Pi_Q)$$
$$\Pi_{1H} > \delta\Pi_Q$$
$$\Pi_Q > 0$$

Given the structure of the problem, this leaves two possibilities: the seller either does or does not sell in the L state:

<u>Case 1: Always sell</u>
The system of equations is:

$$\Pi_L = pL+(1-p)(H-a) + \delta p\Pi_L + \delta(1-p)\Pi_{1H}$$
$$\Pi_Q = p^2L+(1-p^2)H-p(1-p)a + \delta p^2\Pi_L + \delta(1-p)\,\Pi_Q+\delta p(1-p)\,\Pi_{1H}$$
$$\Pi_{1H} = pL+(1-p)(H-a) + \delta p\Pi_L + \delta(1-p)(\,H + \delta\Pi_Q)$$
$$\Pi_{1H} > \delta\Pi_Q$$
$$\Pi_Q > 0$$
$$\Pi_L > \delta\Pi_Q$$

The solution to the system is: $\Pi_Q = \dfrac{ER_Q - \delta p(1-p)(H - ER_L)}{1-\delta}$

and the conditions are equivalent to:

$$\Pi_Q > 0 \Leftrightarrow ER_Q > \frac{\delta(1-p)H}{1+\delta-\delta p}$$
$$\Pi_{1H} > \delta\Pi_Q \Leftrightarrow ER_Q > \frac{\delta(1-p)H}{1+\delta-\delta p}$$
$$\Pi_L > \delta\Pi_Q \Leftrightarrow ER_L > \frac{\delta(1-p)H}{1+\delta-\delta p}$$

Where the last condition implies the earlier two because $ER_Q > ER_L$

Case 2: withhold supply in the low state:

This is the interesting case, in which bargain-hunting persists despite the profitable state being perfectly detectable by the seller.

$\Pi_L = \delta\Pi_Q$
$\Pi_Q = p^2L+(1-p^2)H-p(1-p)a + \delta p^2\Pi_L + \delta(1-p)\Pi_Q+\delta p(1-p)\Pi_{1H}$
$\Pi_{1H} = pL+(1-p)(H-a) + \delta p\Pi_L + \delta(1-p)(H + \delta\Pi_Q)$
$\Pi_{1H} > \delta\Pi_Q$
$\Pi_Q>0$
$\Pi_L> pL+(1-p)(H-a) + \delta p\Pi_L + \delta(1-p)\Pi_{1H}$

The solution is: $\Pi_Q = \dfrac{ER_Q + \delta p(1-p)[ER_L + \delta p(1-p)H]}{(1-\delta)(1+\delta p + \delta^2 p(1-p))}$

With the condition
$\Pi_{1H} > \delta\Pi_Q \Leftrightarrow ER_L > -\delta^2 p(1-p)H$
$\Pi_Q > 0 \Leftrightarrow ER_Q > \dfrac{\delta H(1-\delta p)(1-p)^2}{1+p-\delta p} \Leftrightarrow [p+\delta p-\delta p^2]ER_L > -(1-p)[1+\delta^2 p(1-p)]H$
$\Pi_L > ER_L + \delta\Pi_L + \delta(1-p)\Pi_{1H} \Leftrightarrow ER_L < \dfrac{\delta(1-p)H}{1+\delta-\delta p}$
where $\Pi_{1H}>\delta\Pi_Q$ implies $\Pi_Q>0$

It can be shown that the region of the parameter space satisfying all the conditions is non-empty.

Case 3: Bargain-hunting ceases

Having covered the two bargain-hunting cases, suppose now that $ER_L<-\delta p(1-p)H$, and the seller thus does not sell in state 1H, i.e. does not sell after observing the price of H-a in the state Q. Then, the bidders have no incentive to bargain-hunt, and the game reverts to the game without bargain-hunting. The bidders do not deviate because the seller would witthold supply after noticing a transition to the 1H state.

The game without bargain-hunting is hard to analyze because the seller's ability to learn from prices decreases. Low prices still imply an L state much like the L state above, but a string of High prices leads to beliefs difficult to express in closed form. Please see Appendix 2.3 for details.

One strategy with an immediately available profit is the pulsing strategy, in which the seller receives $p^2L+(1-p^2)H$ every other period. When this revenue is negative, the seller exits the auction market completely. QED

Appendix 2.3: Model without bargain-hunting

Without bargain-hunting, the seller can only observe two possible prices L and H:

Table 2: Prices without bargain-hunting		*New bidder type*	
		L (p)	**H** (1-p)
One old bidder	**L** (p_{1L})	L	H
	H (p_{1H})	L	H
Two old bidders	**L** (p_{2L})	L	H
	H (p_{2H})	H	H

This price-generation changes how much the seller can learn from prices:

Price=L: then the new bidders must have been Low, so:

$$(p_{1L}, p_{1H}, p_{2L}, p_{2H}) \rightarrow \left(\frac{\frac{2}{3}p_{1L} + \frac{1}{2}p_{2L}}{p_{1L} + p_{1H} + p_{2L}}, 0, \frac{\frac{1}{3}p_{1L} + \frac{1}{2}p_{2L} + p_{1H}}{p_{1L} + p_{1H} + p_{2L}}, 0 \right)$$

Price=H: then the only ruled-out possibility is 1L because the winner could not have been of the low type:

$$(p_{1L}, p_{1H}, p_{2L}, p_{2H}) \rightarrow \left(0, \frac{(1-p)\left(p_{1L} + p_{2L} + \frac{2}{3}p_{1H} + \frac{1}{2}p_{2H}\right)}{(1-p) + pp_{2H}}, \frac{pp_{2H}}{(1-p) + pp_{2H}}, \frac{(1-p)\left(\frac{1}{3}p_{1H} + \frac{1}{2}p_{2H}\right)}{(1-p) + pp_{2H}} \right)$$

Unlike the learning of the bargain-hunting model, this learning process visits infinitely many payoff-relevant states. Moreover, the transition probabilities depend on the knowledge in progressively more and more complicated manner. This is because along the "prices=(H,H,H,...)" branch of the game, the probability of yet another price=H is $(1\text{-}p)+pp_{2H}$ where p_{2H} evolves non-monotonically according to the above learning equation. Therefore a simple steady-state analysis on an infinite-horizon model is not readily available.

Appendix 2.4: Model with independent bidders

Always sell:

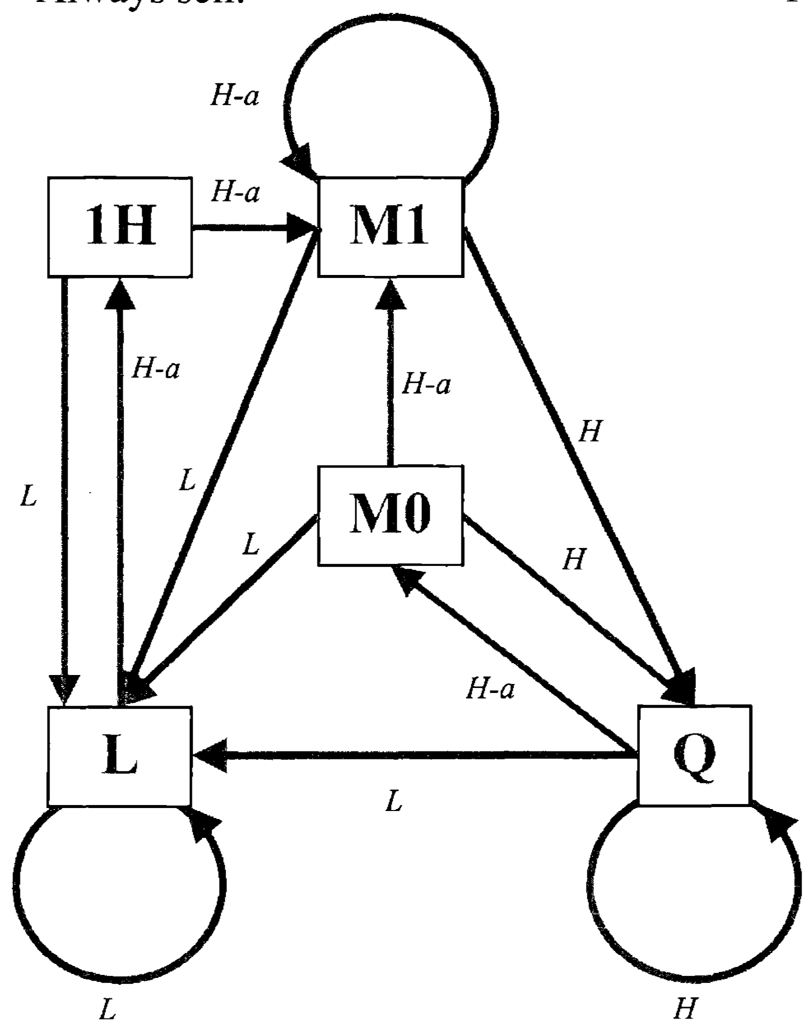

Not sell in L state:

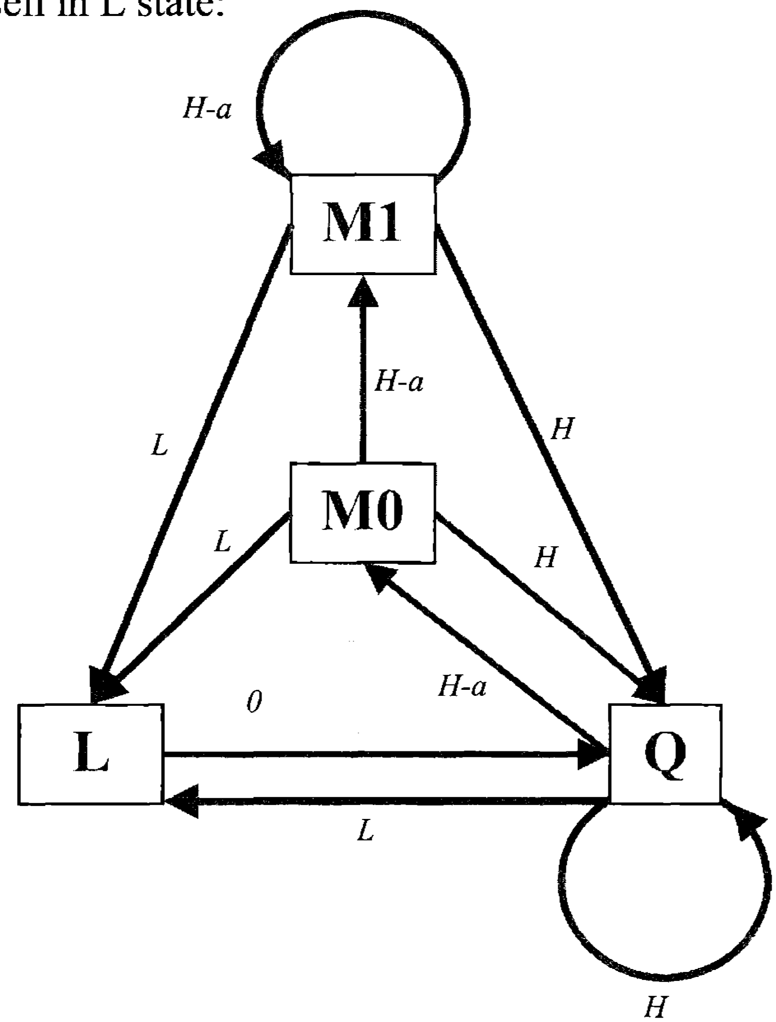

References

Ariely, D. and Simonson, I. (2003): "Buying, Bidding, Playing, of Competing? Value Assessment and Decision Dynamics in Online Auctions", Journal of Consumer Psychology, vol. 13, forthcoming.

Arora, A., Xu, H., Padman, R. and Vogt, W. (2002): "Optimal Bidding in Sequential Online Auctions", working paper, Heinz School of Public Policy Management, Carnegie Mellon University.

Ausubel, L. and Cramton, P. (1998): "Demand Reduction and Inefficiency in Multi-Unit Auctions", working paper.

Bajari P. & Hortacsu A. (2001): "Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions", working paper.

Besanko, D. and Winston, W. (1990) : "Optimal Price Skimming by a Monopolist Facing Rational Consumers", Management Science, vol. 36, No. 5, 555-567.

Coase, R. H., (1972) "Durability and Monopoly," Journal of Law and Economics, vol. 15, 143-149.

Deltas, G. and Kosmopoulou, G. (2002) : "Bidding in Sequential Auctions: 'Catalogue' vs. 'Order-of-Sale' Effects" forthcoming in Economic Journal.

Dholakia, U. and Soltysinski, K. (2001): "Coveted or Overlooked? The Psychology of Bidding for Comparable Listings in Digital Auctions", Marketing Letters, vol. 12:3, 225-237.

Engelbrecht-Wiggans R. (1994): "Sequential Auctions of Stochastically Equivalent Objects", Economic Letters, vol. 44, 87-90.

Engelbrecht-Wiggans R. (1993) : "Optimal Auctions Revisited", Games and Economic Behavior, vol. 5, 227-239.

Gale, D. and Hausch, D. (1994): "Bottom-fishing and declining prices in sequential auctions" Games and Economic Behavior, vol. 7, 318-331.

Greene, "Econometric Analysis", Prentice Hall, 1993

Hausch D. (1986): "Multi-Object Auctions: Sequential vs. Simultaneous Sales", Management Science, vol. 32, No. 12, 1599-1610.

Klemperer, P. (1999): "Auction Theory: A Guide to the Literature", Journal of Economic Surveys, vol. 13(3), 227-286.

Krishna, V. and Rosenthal, R. : "Simultaneous Auctions with Synergies", Games and Economic Behavior, vol. 17, 1-31, 1996.

McAfee, R.P. and Vincent, D. (1997): "Sequentially Optimal Auctions", Games and Economic Behavior, vol. 18(2), 246-276.

Milgrom, P. and Weber, R. (1982a): "A Theory of Auctions and Competitive Bidding", Econometrica, vol. 50, 1089-1122.

Milgrom, P. and Weber, R. (1982b): "A Theory of Auctions and Competitive Bidding II", working paper on multi-unit auctions.

Myerson, R (1981): "Optimal Auction Design", Mathematics of Operations Research, No. 6, 1981, 58-73.

Ortega-Reichert, A. (1968): "Models for competitive bidding under uncertainty", Department of Operations Research technical report 8, Stanford University.

Roth A., & Ockenfels A.(2002): "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet", American Economic Review, forthcoming.

Schauder,J (1930): "Der Fixpunktsatz in Funktionalräumen", Studia Mathematica 2, 171-180.

Vohra, R. and de Vries, S. (2003): "Combinatorial Auctions: A Review", INFORMS Journal on Computing, Vol. 15, forthcoming.

Vulcano, G. and van Ryzin, G. and Maglaras, C. (2002): "Optimal dynamic auctions for revenue management", Management Science, vol. 48 (11), 1388-1407.

Wang, R. (1993): "Auctions versus Posted-Price Selling", American Economic Review, vol. 83 (4), 838-851.

Weber, R. (1983) : "Multi-object auctions", in *Auctions, Bidding and Contracting: Uses and Theory*, eds. Engelbrecht-Wiggans, Shubik and Stark, New York University Press.

Znidarsic, J : "History of the Auction", http://www.auctionwatch.com/awdaily/features/history/, 1999.