NBER WORKING PAPER SERIES

ONLINE AUCTIONS

Axel Ockenfels
David Reiley
Abdolkarim Sadrieh

Working Paper 12785
http://www.nber.org/papers/w12785

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2006

This paper has been prepared as a draft chapter for the Handbook of Economics and Information Systems, Terrence Hendershott, ed., Elsevier Science, 2006. We thank David Caballero and Jason McCoy for their research assistance. Axel Ockenfels gratefully acknowledges support from the German Science Foundation (DFG). David Reiley gratefully acknowledges support from NSF grant SES-0094800. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Online Auctions
Axel Ockenfels, David Reiley, and Abdolkarim Sadrieh
NBER Working Paper No. 12785
December 2006
JEL No. D44

ABSTRACT

The economic literature on online auctions is rapidly growing because of the enormous amount of freely available field data. Moreover, numerous innovations in auction-design features on platforms such as eBay have created excellent research opportunities. In this article, we survey the theoretical, empirical, and experimental research on bidder strategies (including the timing of bids and winner's-curse effects) and seller strategies (including reserve-price policies and the use of buy-now options) in online auctions, as well as some of the literature dealing with online-auction design (including stopping rules and multi-object pricing rules).

Axel Ockenfels
Staatswissenschaftliches Seminar, Universität zu Kö
Albertus-Magnus-Platz
D-50923 Köln
ockenfels@uni-koeln.de

David Reiley
Department of Economics
University of Arizona
401CC McClelland Hall
Tucson, AZ 85721
and NBER
reiley@eller.arizona.edu

Abdolkarim Sadrieh
Faculty of Economics and Management
University of Magdeburg
Postbox 4120, 39016 Magdeburg, GERMANY
sadrieh@ww.uni-magdeburg.de

Online Auctions

Axel Ockenfels, David Reiley, and Abdolkarim Sadrieh*

Abstract

The economic literature on online auctions is rapidly growing because of the enormous amount of freely available field data. Moreover, numerous innovations in auction-design features on platforms such as eBay have created excellent research opportunities. In this article, we survey the theoretical, empirical, and experimental research on bidder strategies (including the timing of bids and winner's-curse effects) and seller strategies (including reserve-price policies and the use of buy-now options) in online auctions, as well as some of the literature dealing with online-auction design (including stopping rules and multi-object pricing rules).

Section 1: Why do information systems make auctions (even) more popular?

Long before any electronic information systems were in place, people used auctions to trade all kinds of goods and services. In his comprehensive overview of the history of auctions, Cassady (1967) reports auctions of items of almost any size, from jewels and spices to ships and provinces. The range of services that have been auctioned is also enormous, including anything from a dance on a local church festivity to the lifetime work force of a slave. While being widespread, however, auctions were not the most common way of trading because the costs of conducting and participating in an auction were typically too high for the every-day trade of common goods. Evidently, the use of auctions is subject to the trade-off between the advantage of *price discovery* (i.e. discovering the highest valuation bidder) and the disadvantage of having high transaction costs (i.e. the costs of finding a buyer and negotiating a sale). Because of this,

---

*This paper has been prepared as a chapter for the *Handbook of Economics and Information Systems*, Terrence Hendershott, ed., Elsevier Science, 2006.

Ockenfels: Department of Economics, University of Cologne, Albertus Magnus Platz, D-50923 Cologne, Germany; ockenfels@uni-koeln.de (http://ockenfels.uni-koeln.de/ao).

Reiley: Department of Economics, University of Arizona, Tucson, Arizona, USA; reiley@eller.arizona.edu.

Sadrieh: Faculty of Economics and Management, University of Magdeburg, Postbox 4120, 39016 Magdeburg, Germany; sadrieh@ww.uni-magdeburg.de (http://ww.uni-magdeburg.de/e-business).

We thank David Caballero and Jason McCoy for their research assistance. Axel Ockenfels gratefully acknowledges support from the German Science Foundation (DFG). David Reiley gratefully acknowledges support from NSF grant SES-0094800.

auctions are most valuable when the party running the auction (both in buy and sell auctions) is very uncertain of the item's market value and, thus, receives a considerable advantage from the price discovery through the auction.

With the emergence and spread of electronic information systems, both the advantages and disadvantages have changed radically. On one hand, the transaction costs associated with conducting and participating in an auction have decreased so substantially, that auctions seem worthwhile even when the expected advantage of detecting the true market value of the item is relatively low. On the other hand, the expected advantage of price discovery has increased sharply, because many more potential auction participants can virtually meet in an online auction house than at any physical location. The effect is magnified by the fact that software agents in virtual auction houses enable participants to interact "on time" without having to be present in person. The time and space emancipation provided by electronic information systems has increased the expected number of potential participants at an auction, making it more likely for the auctioneer to meet participants with especially high valuations.

Since transaction costs have decreased and potential gains of trade have increased, it comes as no surprise that the volume of online auction trades has exploded ever since the service was first offered online. Meanwhile, there is no doubt that online auctions constitute a major market with growing significance for the global economy.

Not only size matters, for in addition to the significance of trade volume, online auctions have also taken a central role in market research and innovation. No other economic institution has induced as much innovation and created as many research opportunities as online auctions have. New features concerning the design of online auctions are proposed and discussed almost on a daily basis. At the same time, the enormous amount of market data that is generated in online auctions and recorded electronically has enabled researchers to empirically address many issues that previously were not researchable. Finally, the open access to online auctions has also opened a door for field experimentation that enables the research of controlled variations in a running market.

In this chapter, we provide an overview of some of the theoretical, empirical, and experimental research on online auctions. The literature in this field is expanding so quickly that comprehensive coverage is impossible. Nevertheless, we have tried to cover all major aspects of the research in the field with the one notable exception being the work on reputation in online auctions. The reason for leaving out reputation in this chapter is that an entire other chapter of this book deals extensively with the issue of online reputation.

The rest of this chapter is organized as follows. In section 2, some of the foundations of auction theory are presented for the simplest case of single-object auctions. The theoretical results are compared to experimental findings and related to empirical observations in online auctions. In section 3, we discuss theoretical and empirical aspects of using public and secret reserve prices in online auctions. Furthermore, we present the research on the relationship between alternative reserve price instruments, such as minimum bids and shill bids. In section 4, we present the research on late and incremental bidding in online auctions. This research deals with the fact that in some online auctions the decisive bids are submitted within the last few moments. An interesting aspect of this strand of research is that it is entirely motivated by phenomena that were first discovered in online auctions. In section 5, we present the research on the buy-now option, which is widespread in online auctions, but rarely observed in classical auctions. The buy-now option creates a hybrid market situation that allows bidders to choose between normal bidding or accepting a posted sale price. While the buy-now option represents an outside option provided by the auction seller, in section 6, we examine parallel auctions and other seller provided options. In section 7, we present the research on multi-object auctions that are both theoretically and empirically more difficult to deal with than the single-object auctions, and in section 8, we conclude the chapter with some general remarks on the design of online auctions.

Section 2: Single-object auctions: theory and experiments

Auction theory has been remarkably influential on the design of electronic market mechanisms. It has also motivated much of the empirical research on auction behavior that we are surveying in this chapter. This section, together with our discussion of multi-object auctions in section 7, reviews some of the central theoretical and experimental results relevant to online auctions.[1]

[1] For more comprehensive and mathematical treatments of auction theory see, e.g., Klemperer (1999), Krishna (2002), Milgrom (2004), and Menezes and Monteiro (2005).

2.1 Standard auction mechanisms and models

An auction elicits information, in the form of bids, from potential buyers regarding their willingness to pay. The outcome – who wins and pays how much – is then determined based on this information. In a single-object auction, one indivisible object is for sale. There are four single-object auction types, which are widely used and analyzed both in theory and practice: the ascending-price auction (sometimes called English auction), the descending-price auction (sometimes called Dutch auction),[2] the first-price sealed-bid auction, and the second-price sealed-bid auction (sometimes called Vickrey auction).

The *ascending-price* auction is probably the best-known auction procedure: the price is raised until only one bidder remains. This remaining bidder wins the object at the price at which the strongest competitor dropped out. There are many ways to run ascending-price auctions: having the seller announce prices, the bidders announce prices, or the price continuously rising on a 'price clock'. In the latter version, which is the one we will refer to when we speak of ascending-price auctions, bidders can quit the auction at any price and observe other bidders quitting. Because the price clock determines the price path, there is no possibility for bidders to speed up or slow down the auction process, or to employ 'jump bids' as a signaling device.

The *descending-price* auction works in the opposite way: the auction starts at a high price, which a price clock then lowers. The first bidder to call out his acceptance of the displayed price immediately stops the clock. This bidder wins the object and pays the price at which the clock stopped. Note that while in the ascending-price auction the winner pays a price determined by his strongest competitor (the price on the clock when the second-to-last bidder exits), the winner in the descending-price auction determines the final price (the price on the clock which he was the first to accept).

In the *first-price sealed-bid* auction, bidders independently submit a single bid, without seeing the others' bids. There is no open, dynamic bidding. The bidder who submits the highest bid wins and pays a price equal to his bid. In the *second-price sealed-bid* auction, again the bidder who submits the highest bid wins, but here he pays a price equal to the second-highest bid.

[2] Outside economics, the term "Dutch auction" is often used for different formats like for eBay's ascending-bid multi-unit auction or, in investment banking, for uniform-price sealed-bid auctions such as in the context of the Google IPO.

In addition to the four types of auctions, there are two standard models of how bidders value an item: the private-value and the common-value model. In the *private-value* model, each bidder knows the value (his maximum willingness to pay) that he assigns to the object, but different bidders may have different values. For auctions to be a meaningful selling mechanism, the seller does not know the values of the potential buyers.[3] Furthermore, there is typically asymmetric information among bidders: one's value is private information to oneself.[4] Provided that there are no entry fees or other costs involved in bidding, the auction winner's net gain is his value of the object minus the final auction price. The losers' net gain is zero since they neither get nor pay anything.

In the *common-value* model, the value of the object is the same to all bidders, but bidders have different information about what is the actual value. For example, the 'true' value of an antique statue may be very similar to all bidders, but bidders may get different 'signals' about whether the statue is genuine or a fake. In such situations, bidders typically want to change their estimates of the value when they learn about the competitors' signals. In the private value model, on the other hand, bidders' values are unaffected by learning the competitors' information. There are also more general modeling approaches, encompassing both special cases of private-value and common-value. In these models, each bidder gets a private signal, and the value to the bidder is a function of all signals.

2.2 Bidding behavior and auction outcomes in theory

We begin with the private-value model. In the *ascending-price* auction, it is a 'dominant strategy' to stay in the auction until the price reaches the bidder's value. That is, no other strategy may yield a strictly higher expected payoff regardless of the competitors' strategies. It follows that the bidder with the highest value wins (the auction outcome is efficient). The price paid by the winner (the auction revenue) equals the value of the strongest competitor, which is the second-highest value of all bidders.

[3] If the seller knew bidders' values, he would simply make a take-it-or-leave-it offer to the bidder with the highest value.

[4] Thus, the appropriate equilibrium concept is Bayesian Nash-equilibrium, in which each bidder maximizes his expected payoff given the competitors' strategies (which are functions of the respective bidder's information) and his beliefs about the other bidders' information.

In the *second-price sealed-bid* auction, a *losing* bid will determine the price. Thus, a bid only affects whether the bidder wins or not, but not his payoff. A bid equal to bidder's value makes sure that the bidder wins if and only if the price is below his value, regardless of the strategies employed by the competitors. Thus, bidding one's value is a dominant strategy (as first observed by Vickrey (1961)). It follows that, as in the ascending-price auction, the bidder with the highest value wins at a price equal to the value of the strongest competitor.

The *first-price sealed-bid* and the *descending-price* auction formats are strategically equivalent. In the descending-price auction, each bidder decides on the price at which he plans to stop the auction. This plan cannot be conditioned on other bidders' bids, because the first bid immediately ends the auction. Similarly, in the first-price sealed-bid auction a bidder must submit a bid without knowing the competitors' bids. Furthermore, the winner pays a price equal to his bid, thus the outcomes of both auctions are equivalent.

In contrast, bidding in the first-price sealed-bid and the descending-price auctions is more difficult than bidding in the other auctions. The only way to make a strictly positive gain is to 'shade the bid', that is to bid less than one's value. For example, in the simple case when the bidders know the values of their competitors, the bidder with the highest value will bid just above the second-highest value. This ensures that he will win at the lowest possible price.[5] If values are private information, however, each bidder faces a risk-return trade-off. The equilibrium strategy then depends on what bidders believe about other bidders' values, as well as their own risk attitudes. The more competition he expects from other bidders or the more risk-averse is a bidder, the higher the optimal bid.

To be more specific, let us assume that bidders are risk-neutral and that values are independent draws from the same distribution, which is the so-called *symmetric independent private-values model*. Then, in equilibrium, each bidder will bid his expectation of the value of his strongest competitor conditional on having the highest value (only in this case is the bid payoff-relevant). As a result, the bidder with the highest value wins (the auction outcome is efficient). The auction revenue is, on average, equal to the second highest value. That is, under our assumptions, the descending-price and first-price sealed-bid auctions yield the same expected revenue as the ascending-price and second-price sealed-bid auctions. Bidders adjust their

[5] If he bids below the second-highest value, the best response of the strongest competitor would be to outbid him.

behavior to changes in the auction rules in a way such that winners do not pay more than what they need to in order to win: the value of the strongest competitor. This is the famous *Revenue Equivalence Theorem* by Vickrey (1961) that has later been generalized by Myerson (1981) and Riley and Samuelson (1981), among others.[6]

The Revenue Equivalence Theorem not only holds for the private-value model, but also for common-value models, if the signals are independent. In auctions with common-value components, however, bidding is more complicated because bidders must take into account that they run the risk of suffering from the *winner's curse*. Each bidder must recognize that (in symmetric equilibria) she wins only if she has the highest estimate of the value. In this sense, winning the auction is bad news: it implies that all other bidders have information that led them to bid more cautiously, so that the winner would probably have revised her value downwards if she had access to competitors' information. The winner's curse refers to the fact that winners may not anticipate this bad news that comes with victory. They run the danger of systematically paying more, on average, than the actual value. Of course, this cannot happen in equilibrium with rational bidding, where bidders would adjust their bids downwards.

Outside the controlled laboratory and field experiments of the sort we survey in the next subsection, the theoretical results described above typically do not directly apply to online auctions, because the circumstances often differ from those assumed in the theory. Bidders may be neither rational nor risk-neutral. They may collude, endogenously enter the auction, or they may be asymmetric with respect to value distributions. Sellers may employ reservation prices or impose participation costs. The same or competing auctioneers may simultaneously or sequentially sell substitutes or complements of the object. There might also be technical risks involved in electronic bidding, and anonymous and geographically dispersed interaction in online auctions may create moral hazard and adverse selection issues. We will address many of these complicating aspects later in this chapter, but first, the next subsection deals with empirical tests of the simple theory.

---

[6] More formally, one version of the Revenue Equivalence Theorem states that if each bidder has a privately known signal (in the private-value model: his private value), independently drawn from a common strictly increasing distribution, then any efficient auction mechanism in which any bidder with the lowest-feasible signal expects zero surplus, yields the same expected revenue.

2.3 Bidding behavior in controlled laboratory and field experiments

Auction experiments in the laboratory and in the field serve as a test of auction theory (Kagel 1995), as an empirical foundation of new approaches in behavioral game theory and other disciplines concerned with economic decisionmaking (Camerer 2003), as a test-bed for alternative auction rules (Roth 2002), and as a teaching tool (Asker et al. 2004). Auction experiments have been conducted in highly controlled laboratory settings to reproduce as accurately as possible the environmental assumptions of a given theoretical model. Particularly in the online-auction era, field experiments can also take place in more natural environments, increasing external validity and decreasing the amount of control by the experimenter. In this section, we concentrate on experimental tests of simple auction theory, focusing on those laboratory and field experiments most relevant to online auctions.

While individual values for the auctioned object are typically unobservable in the field, they can be controlled in experimental settings with the help of the so-called induced-value method (Smith 1976). The trick is to sell money. In a laboratory experiment, bidders compete to win a fictitious good. The bidder who wins the good may subsequently redeem it with the experimenter for a specified amount of cash. This redemption value is typically different for each bidder. So, for example, a bidder with a value of $30 who wins the auction at a price of $21 will earn a cash payment of $9 from the experimenter. By inducing these values and giving information to the bidders about the probability distribution of other subjects' values, the experimenter may impose the assumptions of a given theoretical model.

Laboratory experiments with induced private values demonstrate that bidders tend to bid up to their values in ascending-price auctions, in agreement with theoretical predictions. However, in first-price sealed-bid auctions, bids are typically higher than predicted given the assumptions of the Revenue Equivalence Theorem. This overbidding is a robust effect observed in numerous experiments. Bidder risk aversion was the earliest proposed theoretical explanation for this behavior, but this theory also generated quite a bit of skepticism; see the comprehensive discussion in the survey by Kagel (1995). More recent studies propose explanations based on emotions and bounded rationality. For example, the theoretical papers by Engelbrecht-Wiggans (1989) and Morgan et al. (2003) and the experimental papers by Ockenfels and Selten (2005) and

Engelbrecht-Wiggans and Katok (2005) argue that overbidding is consistent with concerns for relative standing, spite or regret.[7] Kirchkamp and Reis (2004), on the other hand, provide experimental evidence suggesting that overbidding is an artifact of laboratory testing, which often allowed over- but not underbidding on the whole value range.

Laboratory experiments also found that, contradicting the Revenue Equivalence Theorem, open auctions generally raise less revenue but are more efficient than the equivalent sealed-bid auctions (Kagel 1995; see also Engelmann and Grimm 2004 for an analogous result in multi-unit auctions). In particular, descending-price auctions typically raise less revenue than first-price sealed-bid auctions, and bidders in second-price auctions often overbid with respect to their dominant strategy and rarely underbid (Kagel and Levin 1993, among others). However, recent experimental studies in environments that are closer to 'naturally occurring' online auction environments sometimes seem to qualify these findings.

In the first online-auction experiment, Lucking-Reiley (1999) auctioned off collectible trading cards over the Internet. By going into the field, he sacrificed some control – e.g., he did not induce values, he allowed for endogenous entry etc. – in order to assess the predictive power of the theory in a 'real-world' auction. He found that descending-price auctions earn approximately 30 percent more revenue than first-price sealed-bid auctions, which is inconsistent with both the Revenue Equivalence theorem and previous laboratory results. He could not find a significant difference between the ascending-price and the second-price sealed-bid formats,

Other experiments exploit the fact that many online auction platforms, such as eBay, operate much like second-price auctions. EBay asks the bidders to submit maximum bids (called "proxy bids") and explains that "*eBay will bid incrementally on your behalf up to your maximum bid, which is kept secret from other eBay users*." That is, once a bidder submits his (proxy) bid, eBay displays the currently winning bid as the minimum increment above the previous high proxy bid.[8] At the end of the auction, the bidder who submitted the highest bid wins the auctioned item and pays a price equal to the second-highest bid plus the minimum increment..[9]

[7] Recent social comparison models such as Bolton and Ockenfels (2000) and Fehr and Schmidt (1999) proved to be quite successful in capturing 'anomalous' behavioral patterns in a wide range of economic situations.

[8] While proxy bidding became widespread with the advent of eBay, Lucking-Reiley (2000b) documents similar rules being used by auctioneers who wished to allow absentee bids, dating back at least to the 1800s.

[9] If the second-highest bid plus one increment exceeds the highest bid, then the highest bidder pays his own bid.

In the analysis that follows, we shall assume the minimum increment amount to be negligibly small.

To understand the connection between the single-unit eBay auction and the second-price sealed-bid auction, assume for the moment that nobody learns about the proxy bids of other bidders until the auction is over. Then, in fact, eBay becomes a second-price sealed-bid auction, in which each bidder has a dominant strategy to bid his value. EBay explains the economics of second price auctions to their bidders along these lines, and extends the conclusion to its own auctions, in which bids are processed as they come in: "*eBay always recommends bidding the absolute maximum that one is willing to pay for an item (...)*." Ockenfels and Roth (forthcoming) support this view based on a game-theoretic analysis of a continuous-time model of eBay's single-object auction. They show that, while there is no dominant strategy in eBay's open second-price format, strategies that involve bidding above one's value are dominated, and that bidders 'sooner or later' will always bid their value.[10]

Garratt, Walker and Wooders (2004) investigated bidding behavior of eBay buyers and eBay sellers, experienced with eBay's second-price rule, using induced values in a second-price sealed-bid auction experiment conducted on the Internet. While even highly experienced eBay bidders tend to not bid exactly equal to their values, there was no greater tendency to overbid than to underbid as previously observed in laboratory experiments. Furthermore, Garratt et al. found that, on average, subjects with selling experience on eBay bid significantly less than subjects with only buying experience on eBay. One important implication of this study is that background and experience with online auctions seems to affect bidding behavior.

Ariely et al. (forthcoming) investigated eBay bidding behavior in the laboratory by programming an experimental auction that emulated eBay's central auction rules as described above. In one study, they compared the performance of eBay's open second-price auction format with an analogous second-price sealed-bid format. For inexperienced bidders, they reported that the median sealed bids were substantially lower than in the open auction. Consistent with previous laboratory studies, the open format results in more efficient outcomes, but inconsistent with previous results, the open format also yields higher revenues. One reason for the better

---

[10] To be slightly more specific, this holds for all Bayesian equilibria in undominated strategies of their private-value auction model.

performance of the open format is that learning in the sealed-bid auctions only takes place across auctions, while learning in the dynamic auctions also takes place within auctions. For example, a bidder who imagines that he can win with a low bid does not learn that he is mistaken in a sealed-bid auction until after the auction is over. However, in open auctions, he can revise his bid as soon as someone outbids him. For experienced bidders, Ariely et al. reported median final bids in both the sealed-bid and the open auctions that converge to the equilibrium prediction: 100 percent of values.

Ockenfels (2005) studied bidder behavior in a one-shot field experiment, using eBay as the experimental platform. He invited eBay bidders through posters and email to register for online, eBay experiments; the only requirement was a valid eBay account. Bidders were privately informed about their induced values. All communication and data collection were automated and via the Internet. Due to the second-price rule employed by eBay, final bids should theoretically be equal to induced values. Ockenfels found that, on average, the losers' final bid (eBay's interface does not reveal the winners' maximum bids) were indeed close to the induced values.

Regarding common-value auctions, laboratory experiments consistently show that inexperienced bidders are subject to the winner's curse: on average, winners typically overpay (Kagel and Levin 2002). Learning to avoid the winner's curse appears difficult. Ball et al. (1991) explored the winner's curse in a bilateral bargaining game with asymmetric information, and found virtually no adjustment to the winner's curse over twenty trials (see also Grosskopf et al. 2003). Kagel and Levin (1986) observed that experienced bidders in first-price, common-value auctions can overcome the winner's curse only with relatively few bidders, but succumb again with larger number of bidders. Observe that a larger number of bidders heightens the winner's curse problem – winning is worse news when there are more competitors who think the object is not worth the final price. So, in theory, more bidders require further shading of bids, while, in fact, the subjects in the experiment by Kagel and Levin (1986) tended to bid more aggressively in larger bidding groups. Kagel et al. (1995) replicated this finding in laboratory, second-price, common-value auctions. Bidders again failed to respond in the right direction to more competitors.

Given the laboratory evidence, it is reasonable to expect that the winner's curse also plays a role in online auctions, because many online auctions have common-value components. For one, online auctions often make it too costly for buyers to inspect in person the object being sold, so that an assessment of the 'true' condition of the object can be difficult and may vary across bidders, depending on the sources and quality of individual information. Moreover, there is often a risk of being exploited by a fraudulent seller on C2C auction platforms. Because bidders may differ with respect to their assessment of a seller's trustworthiness, winning the auction may imply bad news regarding the seller's intentions. Finally, online auctions often attract large numbers of potential bidders, which further amplifies the winner's curse; winning an object on eBay means, in the extreme, that thousands of other users do not think that the object is worth its final price.

A couple of papers looked for winner's curse effects on eBay. While direct evidence from field experiments supports the laboratory findings that winner's curse effects are a real and robust phenomenon,[11] indirect statistical analyses of eBay data also support the view that bidders, to some extent, respond in a strategic way to take the risk of overpayment into account.

Jin and Kato (2004, 2005) conducted a field experiment searching for the winner's curse on eBay. They bid on eBay auctions for upgraded, baseball cards, and then let a professional grading service evaluate the cards. They found that claims of high value result in more fraud (i.e., default or sending counterfeits) and no better card quality. However, eBay buyers were willing to pay 23 to 54 percent more for cards that claimed quality of mint or better. Overall, Jin and Kato concluded that some buyers have fallen prey to the winner's curse by having underestimated the probability of fraud (see also Bajari and Hortacsu 2004 for a discussion of these findings). A more indirect test by Bajari and Hortacsu (2003) relied on analyzing strategic bidding within a common-value auction model of eBay. The model is a second-price, sealed-bid auction with an uncertain number of bidders. The sealed-bid aspect of the model implies that no bidder can learn about the value from the bids of others.[12] Using field data of collectible coin auctions and applying various statistical instruments to account for the endogeneity of the number of bidders, the authors found that bids decline with the number of competing bidders. This is in line with theory but inconsistent with laboratory research. Regarding experience effects, Bajari and Hortacsu found that experienced bidders are slightly more cautious. On the other hand, Ow and

[11] A related but distinct phenomenon is price dispersion: bidders often pay a higher price for an object than prices for an identical object offered at the same time by a different seller; see, e.g., Ariely and Simonsohn (2003) and the survey by Morgan, Baye and Scholton (2005) in this Handbook.

[12] The authors justify this modelling approach by noting that many bids are submitted in the last minute of the auctions.

Wood (2004) reported, also based on a field study of rare-coin eBay auctions, that more experience leads to more aggressive bidding. However, this is not necessarily inconsistent with the winner's curse story put forward by Bajari and Hortacsu, since Ow and Wood argue that the winner's curse effect is partly overlaid by an "initial lack of institution-based trust" that decreases with experience. Other evidence in favor of rational reactions to winner's curse effects comes from survey information on completed eBay auctions by Yin (2005). She showed that the winning bid in a sample of eBay auctions is negatively correlated with the variance of the self-reported, ex-ante estimates of the objects' values.[13] That is, the more certain a bidder is of the item's value, the more aggressively they bid.

The data surveyed here suggest that auction theory and laboratory results are sometimes, but not always, a good predictor of online auction behavior. For instance, auction theory has difficulties capturing overbidding in first-price auctions and overpaying in common-value environments. Field behavior seems to differ from laboratory behavior when it comes to overbidding in second-price auctions and to experience and number effects in common-value environments. Some causes for these discrepancies suggest themselves: lack of bidder experience (Garrat et al. 1994); small decision costs and stakes (List and Lucking-Reiley 2002); uncontrolled institutional differences, self-selection and subject-pool effects; and presentation effects.[14] However, to date, there is too little research on the external validity of theoretical and experimental auction research. One of the exceptions is Bolton and Ockenfels (2006), who had 'real' experienced eBay traders interact in the 'real' eBay market, played out by eBay's rules, and they do so with no sacrifice of control over theoretical predictions. They are able to partly link the behavioral patterns observed in their controlled experiments on eBay both to behavior observed in fully controlled laboratory experiments and to behavior observed in the fully uncontrolled eBay market. The data confirm laboratory results (predicted by social preference theory), that buyer competition yields highly competitive results, while bilateral trading environments yield more equitable outcomes, which are inconsistent with standard auction and

---

[13] High variances occur in auctions with poorly designed web pages or where the object had ambiguous characteristics.

[14] For the latter, observe for instance that eBay's 'proxy bidding' explanation of second-price auctions ("... *think of the bidding system standing in for you as a bidder at a live auction"*) appears much more understandable than a typical explanation used in a laboratory experiment (*"You pay the amount of the second-highest bid -- i.e., the highest bid placed by any other participant in the auction."*)

bargaining theory. At the same time, the data show that a priori experience with eBay affects the sellers' trading strategies in the field experiment. It appears likely to us that, because online auctions allow "losing control in a controlled way" (Bolton and Ockenfels 2006), they will play an important role in investigating how laboratory and field behavior are linked to each other.

Section 3: Reserve prices, minimum bids, and shill bids

Superficially, it may seem that all activity in an auction comes only from bidders submitting offers, while the seller running the auction simply waits and hopes that the auction will be profitable. However, in addition to the initial choice of auction mechanism, the seller can also choose a *reserve price* that prevents low-revenue sales and stimulates competition. Strictly speaking, a reserve price defines the minimum amount that a seller will accept for the auction to end with a sale.[15]

The most straightforward behavior for the seller is to set a reserve price equal to her true willingness to accept and announce it at the outset of the auction. In this case, the reserve price would serve to enhance efficiency, preventing the item from being sold if none of the buyers values it as highly as the seller. However, the seller may also make strategic use of the reserve price by setting it higher than her willingness to accept. The theoretical results discussed in the next subsection will show that optimal reserve prices are often – but not always – set strictly higher than the seller's willingness to accept.

In most online auctions, the seller may make a strategic choice not only of the amount of the reserve price, but also whether the reserve price should be secret or public and, if made public, at what point in the auction. Finally, though this violates the formal rules of the auction game at eBay and most other online auction sites[16], the seller may effectively camouflage and dynamically adjust the reserve price during the auction by using *shill bids*, or bids covertly placed by the seller or by the seller's confederates to artificially inflate the final sale price. Clearly, any

[15] For convenience, we restrict our attention to the case where the auction is held by the seller, but the same discussion applies equally well (with obvious adjustments) to procurement auctions held by a buyer.

[16] Most online auctions explicitly forbid shill bidding on their sites. eBay, for example, writes: "Shill bidding is the deliberate use of secondary registrations, aliases, family members, friends, or associates to artificially drive up the bid price of an item. (This is also known as "bid padding.") Shill bidding is strictly forbidden on eBay. EBay members found shill bidding may receive warnings or indefinite suspensions of their accounts." And: "Shill Bidding undermines trust in the marketplace since prices are being artificially manipulated. Furthermore, it is illegal in many locations throughout the world. To ensure that buyers' and sellers' trust in eBay's marketplace is appropriately placed, Shill Bidding is not permitted."

of these strategic options (or combinations thereof) may be used by the seller to increase the expected revenue from the auction.

The most popular type of reserve price used in online auctions is the *minimum bid* (sometimes also called the *opening bid* or *starting bid* in dynamic auctions). A minimum bid defines the lowest bid that a bidder may submit at the outset of an auction. Because it is publicly announced before the auction begins and cannot be adjusted later, a minimum bid represents a static *public reserve price*.

When the seller sets the minimum bid below her valuation, she often combines this strategy either with a *secret reserve price* or with shill bidding by the seller or a colluder. Neither of these reserve-price variants are made public; in fact, shill bidding is a type of fraud. However, both have similar effects to public reserve price: a trade only occurs if the final highest bid is above the secret reserve price or the shill bid. The three instruments differ in their informational and dynamic aspects. Secret reserve prices are fixed with the auctioneer before the auction starts. On eBay and most online auction sites, bidders are informed that an auction has secret reserve, and whether or not it has yet been met by the bidding. (In an interesting contrast, traditional, live auction houses like Christie's and Sotheby's do not inform bidders whether any secret reserve price has yet been exceeded.)

By contrast with secret reserve prices, shill bids – like all bids in dynamic auctions – can be submitted and updated at any time during the course of the auction. Bidders are not informed of the presence of shill bidding, but obviously, they might guess that shill bidding is taking place. These differences in the informational and dynamic features are not only theoretically relevant, but in many countries also have legal consequences. It is important to note, however, that online shill bidding can be well organized (e.g. with a relatively large set of paid colluders, a rotating scheme with peer sellers, or through the use of a false online identity) and hard to detect. The possibility to use any number of anonymous online identities has substantially simplified undetectable shill bidding.

3.1 Theoretical considerations

A public reserve price can increase the revenue of a seller, both in the independent-value setting (Myerson 1981, Riley and Samuelson 1981) and in the affiliated or common-value settings (Milgrom and Weber 1982). In the first case, the seller should choose a reserve price that maximizes his expected income by extracting the expected surplus from the highest bidder. This optimal reserve price will typically be well-above the seller's valuation. Since the seller does not know the realized buyer valuations, the optimal reserve price will sometimes turn out to be higher than the highest bidder's valuation, in which case no trade will occur. Hence, analogous to the case of monopoly pricing, the optimal reserve price raises expected revenues for the seller but leads to inefficiency through reducing the quantity of trade. A number of authors have argued that this inefficiency cannot persist, because it entails incentives for re-negotiations and re-auctioning, unless the item is extremely perishable. The modeling of post-auction resale leads to different conclusions about equilibrium auction outcomes and optimal reserve prices (Haile 2000, Zheng 2002).

In the case of affiliated and common values, where bidders do not know their own values with certainty, the seller's revenue will usually increase with a public reserve price. Since the announcement of the reserve price may improve the bidders' information about the true value of the auctioned item, the bidders can reduce their cautious bid shading and submit higher bids (Milgrom and Weber 1982). At the same time, however, a public reserve price may reduce the amount of information available to the active bidders in an ascending bid auction. Since in this case the bidders with signals below the reserve price cannot participate, their signal information is not revealed in their bids.

The theoretical results mentioned above are based on the assumption of a fixed and known number of bidders, who incur no cost of entry. When the number of bidders is endogenous (i.e. bidders can choose whether or not to participate) and bidders have some cost of entry, it may be advantageous for the seller to set the reserve price no higher than her valuation in order to encourage efficient levels of entry (Samuelson 1985, Engelbrecht-Wiggans 1987, McAfee and McMillan 1987, Levin and Smith 1994).

The theoretical effects of secret reserve prices are also rather mixed. The obvious market-design question is whether the use of a secret reserve price is more beneficial than a public reserve price (minimum bid). Tatsuya Nagareda (2003) models a second-price, sealed-bid auction where the seller can either set a public or a secret reserve price. He finds that no symmetric equilibrium exists in which secret reserve prices increase the expected revenue of the seller. Other researchers, such as Elyakime et al. (1994) analyze an independent-value, first-price auction and conclude that a seller is strictly worse off using a secret reserve price versus a minimum bid.

Not all theoretical models predict a disadvantage to secret reserve pricing. Li and Tan (2000) focus on risk-averse bidders rather than risk-neutral bidders. The authors find that with risk-averse bidders, a secret reserve may increase the seller's revenue in an independent, private-value, first-price auction. On the other hand, in second-price and English auctions, risk preference does not play a role and the seller should be indifferent between a private or public reserve price. The work of Vincent (1995) provides an example where setting a secret reserve price in an English or second-price auction can increase a seller's revenue in an affiliated-values setting. He argues that since a nonzero minimum bid can cause some bidders to avoid the auction entirely, the attending bidders will have less information than in an auction with a secret reserve price, but no minimum bid. As usual, less information on other bidders' signals in an affiliated-values auction leads to more cautious equilibrium bidding and hence lower prices.

With shill bidding, the informational situation is rather different than in the case of secret reserve prices. For one thing, the bidders in such auctions receive no explicit notice that seller is effectively imposing a reserve. Of course, in an institution in which shill bidding is possible, buyers may expect it to happen. In fact, Izmalkov (2004) shows that in an English auction with asymmetric independent private values, there exists an equilibrium with shill bidding that has an equivalent outcome to that of Myerson's (1981) optimal mechanism. The intuition for this result is best described by Graham et al. (1990 and 1996), who show that setting a reserve price dynamically, that is after having observed some bidding, can increase the seller's revenue. The effect is due to the fact that the longer the seller can observe bidding in the auction, the more precise becomes the seller's information buyers' values. In the absence of penalties for shilling, the seller should weakly prefer an adjustable reserve price (e.g. a shill bid) to an ex-ante fixed reserve price.

Similar, more detailed results for the independent-private-value ascending auction are shown by Sinha and Greenleaf (2000) and Wang et al. (2004). For example, Sinha and Greenleaf's discrete-bidding model shows that the advantage to the seller of shilling depends on a number of parameters, including the number of active bidders and their "aggressiveness" as well as the sequence of buyer and shill bids. They find that, in some cases, sellers should optimally refrain from shill bidding and credibly commit to their abstinence. This commitment ensures that bidders shade their bids less than when they fear shill bidding occurs.

Chakraborty and Kosmopoulou (2004) derive a similar result for the case of dynamic shill bidding in a common-value setting with an ascending auction. Their model uses a pool of bidders who are only able to see a binary signal for the value of the good (high or low). They show that as a seller increases his rate of shill bids, while holding bidder behavior constant, he increases his selling price, since common-value bidders infer higher item value from greater participation. However, since bidders guess that the seller will employ shill bidding, they decrease the amount of their bids, which lowers the final sale price. Furthermore, when the seller happens to win his own auction with a shill bid, he must pay a fee to the auctioneer without actually making a sale. If these two negative effects outweigh the potentially higher selling price, the seller would prefer to commit to a policy of no shill bidding. However, the seller has a credibility problem with committing to not shill, because in any auction where bidders do not believe shill bidding is occurring, a seller has a clear incentive to shill in order to increase the final price. Given the seller's lack of credibility, bidders should always believe that shill bidding will occur and lower their bids accordingly.

Finally, some features unique to online auctions make shilling behavior more attractive to the seller. Engelberg and Williams (2005), for example, argue that dynamic shill bidding is strongly supported by eBay's system of "proxy bidding" to approximate a second-price auction. The bidder with the highest bid in an eBay auction is called the "high bidder" and holds the "current bid" that usually is equal to the next highest proxy bid plus a fixed bidding increment. One important exception is that if the proxy bid of the high bidder is not large enough to provide a full minimum increment over the second-highest bid, then the current bid is set to exactly the value of the high bidder's proxy bid. In that event, the second-highest bidder can infer that the high bidder's proxy bid amount has just been reached.

Engelberg and Williams point out that this feature of eBay's rules facilitates a "discover-and-stop" shilling strategy.[17] They observe that most bidders enter bids of whole- or half-dollar amounts, so a particularly effective shilling strategy would place bids with unusual decimal parts, for example, making all shill bids end in 37 cents. Continuously increasing the shill bid by the

---

[17] Shah et al. (2002) describe a similar pattern of "uncovering" behavior, without the special eBay details.

minimum bid increment up to the point in which the current bid is no longer increased by the full amount of the minimum bid increment, allows sellers to squeeze the full value of the item from the highest bidder, while minimizing the chances of winning the auction and receiving no revenue. For example, suppose that the high bidder has submitted a proxy bid of $7.50 and that the minimum bid increment is 50 cents. The seller first submits a shill bid of $5.37, which results in eBay showing the high bid as $5.87. He continues with a series of shill bids increasing in one-dollar steps: $6.37, $7.37. At this point, the system reports a current bid of $7.50 instead of $7.87, revealing the high bidder's maximum bid to be exactly $7.50. At this point, the shill bidder stops, extracting the bidder's full willingness to pay, without risking overbidding the true bidder and failing to sell the item. Engelberg and Williams (2005) conjecture that a bidder's best response to this discover-and-stop shill strategy is one of sniping, or withholding their true bid until the last seconds before the closing of the auction (see Section 4).[18]

Overall, it seems clear that the theoretical predictions concerning the different variants of reserve prices depend on many details. Clearly, all of the reserve-price variants can, under some circumstances, increase the seller's revenue. However, the rational reaction of the buyers often involves increased bid shading. As the sellers become smarter in their effort to uncover bidder's valuation information, bidders should become ever more cautious about revealing information in their bids. Hence, empirical and experimental work is needed to assess the actual economic impact of reserve prices and to compare the effects of different reserve-price strategies.

3.2 Empirical and experimental observations

Many of the theoretical results concerning public and secret reserve prices depend on the details of the models used. The effects of reserve prices are especially sensitive to theoretical assumptions about the information and valuations of bidders. While the empirical and

---

[18] In an environment without proxy bidding, Kirkegaard (2005) discovers a related theoretical result: bidders may wish to submit jump bids as a defense against shilling. By mixing jump bids and normal bids, bidders obscure the information about their willingness-to-pay, making dynamic shilling less effective. Note that proxy bidding precludes jump bidding, because a single bidder can't cause the current bid to jump above the second-highest bid; someone wishing to execute a real jump bid on eBay would have to use two different bidder accounts. Note also that Kirkegaard uses the term *phantom bidding*, which is technically distinguished from shill bidding because the former refers to the auctioneer making up a non-existent (i.e. "phantom") bid, while the latter refers to the seller or a confederate actually placing a bid with the auctioneer.

experimental research we introduce in this subsection provides some help in sorting out and matching the appropriate models and auction situations, we will also see many questions left open for further research.

3.2.1 Entry and revenue effects of public and secret reserve prices

Although details may differ, there are a few predictions that are shared by all theoretical models. These basic predictions are amongst the earliest studied in the field. The first basic hypothesis is that reserve prices (whether public or secret) should lead to a decrease in the number of bids and the number of bidders in an auction. The second hypothesis is that the number of auctions ending without a trade should increase when reserve prices are used. What overall effect these two observations have on average prices (i.e. revenues) depends on the details of the theory used.

An early test of these hypotheses was performed by Reiley (2006). In his field experiment, collectible trading cards from the game "Magic: The Gathering" were sold in first-price, sealed-bid auctions on Internet newsgroups. The size of the minimum bid (public reserve price) was varied systematically as a fraction of each card's book value, or reference price . The main results of the experiment are consistent with the basic hypotheses above: holding all other variables constant, the use of a public reserve price (1) reduces the number of bidders, (2) increases the frequency with which goods go unsold, and (3) increases the revenues received on the goods conditional on their having been sold. Furthermore, bidders clearly exhibit strategic behavior in their reactions to the public reserve prices. High-value bidders, for example, raise their bids above the reserve in anticipation that rival bidders will do the same.

Ariely and Simonson (2003) study eBay auction prices for tickets to the 2000 Rose Bowl (a popular, American, collegiate football game). They found that the minimum bid and the total number of bids have a positive correlation to the price. Unfortunately, the authors do not report the interaction between the minimum bid and the number of bids. According to the first basic hypothesis and given the evidence from the other empirical studies, we should expect that the number of bids will depend on the value of the minimum bid and cannot simply be viewed as an exogenous parameter. In fact, in a follow-up field experiment on VHS, DVD, CD, and Book sales on eBay reported in the same paper, the authors observe that bidder activity and the number of bids submitted were greater for a low minimum bid than for a high one. This clearly indicates that activity measures such as the number of bids should be treated as endogenous parameters in empirical work on auctions.

With their field experiment, Ariely and Simonson (2003) show that there is another exogenous parameter that may affect the market activity level and may interact with the effect of the minimum bid on prices. This exogenous parameter is a measure of the amount of supply by other sellers. When many sellers are offering identical (or very similar) items at the same time,[19] then auctions with both high and low minimum bids end at roughly the same price. Thus, a high degree of supply by other sellers reduces the effect of the public reserve price. By contrast, when there are few other sellers offering the same item, a high minimum bid yields empirically higher auction prices. Note that this effect is not in line with standard auction theory. In a standard auction model, we would expect high seller-side competition to lead to a lower number of bidders per auction. This, in turn, should make the use of a public reserve price more valuable to the seller, because it helps to drive up the price especially in "thin" markets. Hence, standard auction theory would predict the greatest difference between auctions with low and high minimum bids when seller-side competition is high. The authors explain their results by claiming that the low seller-side competition reduces the probability that bidders compare competing auctions and, thus, enables the minimum bid in each auction to be a much more effective psychological "anchor" for the bidding behavior.

The anchoring hypothesis finds some support in the field and laboratory auction experiments reported by Häubl and Popkowski Leszczyc (2003). In their field experiments, they auction identical postage stamps, while systematically varying the minimum bid, and the shipping and handling costs. They find that the total selling price substantially increases with the minimum bid and the fixed cost of shipping. The effect seems especially strong when the true value of the item is harder to estimate. However, an external reference point does not alleviate the effect.

Hossain and Morgan (2006) conduct eBay field experiments for Xbox games and music CDs, systematically varying the shipping costs. They find that for Xbox games, setting a low minimum bid and a high shipping cost yields more revenue than doing the reverse. Buyers do not seem to take the extra shipping cost as much into account as the stated minimum bid in the

[19] Ariely and Simonson term this "high comparability."

auction. They do not discover the same effect for music CDs; when the shipping costs were a substantial fraction of the item's selling price, the bidders took shipping costs into account just as much as minimum bids. These results are consistent with psychological anchoring effects being present but limited in scope.

Anchoring can obviously be effective only with public reserves, not secret ones. Hence, if the anchoring hypothesis is true, we should expect to observe a greater correlation between a public reserve price and the auction outcome, than between a secret reserve price and the auction outcome. Unfortunately, we know of no study so far, that systematically varies the amount of the secret reserve price. There are, however, a number of studies comparing public and secret reserve prices.

In an early empirical study, Bajari and Hortaçsu (2003) examined the effects of minimum bids and of secret reserve prices in all 407 auctions for mint and proof sets of U.S. coins that occurred on eBay during a week in late September of 1998. Only 14% of the observed auctions used a secret reserve, with the average book value of these items being more than 20 times higher than the average of items without a secret reserve. While 84% of the items without a secret reserve sold, only 49.1% of the items with a secret reserve sold. Surprisingly, the average number of bidders for auctions with a secret reserve was substantially higher (5.0) than in the other auctions (2.8). This correlation disappears when the confounding effect of item value is taken into consideration: high-value items generally induce more bidding activity, and are also more likely to have secret reserves. However, these high-activity auctions are also associated with lower bids relative to book values. Overall, the results suggest that a secret reserve has less of an entry-deterring effect than a public reserve, but a secret reserve does have a positive effect on revenue. Hence, the authors suggest that a combination of a low minimum bid and a secret reserve is likely to be the optimal configuration from a seller's point of view, especially in auctions of high-value items.

Dewalley and Ederington (2004) also analyzed eBay auctions to measure the impact of secret reserve prices on bidding strategies and final auction prices. Their data on 5,275 auctions of classical Silver Age comic books was gathered in 2001 and 2002. Unlike Bajari and Hortaçsu (2003), they find clear statistical evidence that the use of a secret reserve reduces the number of active bidders in an auction and, thus, has a negative impact on the seller's revenue. This result is strongly supported and extended by Katkar and Reiley (2005), who report on a field experiment in which they auctioned on eBay 50 matched pairs of Pokémon trading cards. One card of each pair was auctioned with a minimum bid, while the other was auctioned with an equivalent secret reserve price. On average, the secret-reserve auctions return 10% less revenue and are more than 30% less likely to end in a sale.

3.2.2 Auction fever

Auction fever is one of the most frequently discussed issues concerning online auctions.[20] In general, auction fever is thought to be an excited and competitive state-of-mind, in which the thrill of competing against other bidders increases a bidders' willingness to pay in an auction, beyond what the bidder would be willing to pay in a posted-price setting. Since auction fever supposedly derives from the thrill of competition, one might reasonably expect the effect to increase with the number of active bidders. This theory may explain why some auctioneers prefer a low minimum bid, perhaps lower even than the auctioneer's true willingness to accept. The low minimum bid would attract as many bidders as possible, in an attempt to promote auction fever. (In case auction fever is insufficient, shill bidding could prevent the item being sold below the seller's reservation price.)

It is important to note that auction fever (sometimes also referred to as *competitive arousal*, *bidding frenzy*, or *bidding war*) generates a diametrically opposite prediction to both the standard auction-theoretic argument for reserve prices and the anchoring hypothesis described above. While those imply that installing a high public reserve price will increase seller's revenues, auction fever predicts that a low public reserve price will create a more competitive atmosphere, which in turn leads to bidders' arousal, higher bids, and higher auction revenues.

Häubl and Popkowski Leszczyc (2004) run a series of experiments in which they vary the frequency of bid arrivals and the perceived total number of active bidders in order to find evidence for auction fever.[21] They find both parameters to have positive, significant effects on

---

[20] In many of its ads, eBay actually advertises with an image of aroused bidders enjoying the thrill of bidding and the joy of winning.

[21] To be able to manipulate these parameters in a controlled manner, Häubl and Popkowski Leszczyc (2004) let each subject play against bidding robots that are programmed to create each treatment environments. In order not to disturb the emergence of auction fever, the authors mislead their subjects to believe that they are bidding against

revenues, indicating that auction fever may be effectively pushing up prices. No such effect is observed when bid increments and time pressure are varied.

Ku, Malhotra, and Murnighan (2005) explore field and survey data of online and offline auctions to look for evidence for competitive arousal. The survey results seem to provide evidence for auction fever. In addition to the evidence for auction fever, the authors also find evidence of overbidding due to an attachment effect, which is when long bidding durations and other sunk costs intensify the desire to win the auction, leading to increased revenues for the seller. Both effects are also observed in a controlled laboratory experiment, in which the sunk-cost parameter and the number of bidding rivals was varied.

Heyman, Orhun, and Ariely (2004) also examine these two phenomena of competition and attachment. They use the term "opponent effect" to describe the arousal due to competing with others and the term "quasi-endowment" for increased valuation due to having been attached to the item as the high bidder over a long period. In two experiments, one comparing different hypothetical scenarios, and one examining real transactions in the laboratory, they vary the number of rival bids and the duration of quasi-endowment (i.e. time spent as the high bidder). Both an increase in the number of rival bids and an increase of the duration of quasi-endowment have a positive effect on the final price. The authors conclude that sellers may be able to increase revenues by increasing the total auction duration and by lowering the minimum bid in order to induce more "feverish" bidding.

The evidence to date suggests that auction fever may be a real phenomenon. This implies that sellers might be able to increase revenues by setting very low minimum bids in order to increase the number of active bidders; however, that prediction has not yet been tested directly.[22] These studies also report somewhat higher bids from those bidders who were the high bidder for a longer duration (and hence would have more opportunity to become "attached" to the item).

---

other human players. Strictly speaking, it is not clear to what extent subjects from a subject pool that has been exposed to deception in experiments, will actually believe any part of the information that they are given.

[22] Reiley (2006) presents some evidence along these lines in field experiments involving first-price sealed-bid auctions: realized prices were slightly higher with no minimum bid than with a minimum bid at 20% of book value. However, "auction fever" typically refers to ascending-bid auctions where bidders can observe the number of competitors, not to sealed-bid auctions where the number of competitors is uncertain.

3.2.3 Shill bids

Identifying shill bids in field data is an extremely difficult task, even though most online auction sites provide detailed bid-history information for every auction. The most crucial issue is that online identities are easily created and cannot be tracked back to the physical identities without inside information, nor is it easy to prove collusion among legitimate seller identities. This means that, in theory, a seller can create any number of online buyer identities and have them bid on the items auctioned. Obviously, creating online identities is not free of cost: at the very least, the seller will incur opportunity and effort costs. In practice, this means that sellers who shill will have some incentive to try to minimize the number of fake identities they use to cover up the shill bidding. But even with a small number of fake online identities, identifying the connection between a seller and the seller's *confederates* remains a difficult empirical task. Proving such a connection is even more difficult.

Some authors have presented clever approaches to this difficult problem. Kauffman and Wood (2003) gathered data from eBay auctions of rare coins. Their central instrument for the detection of shill bidding consists of the search for "questionable bids," meaning bids that appear to be strictly dominated from the bidder's point of view, but could be rationalized as a seller's shill bid. Kauffman and Wood (2003) consider the following criteria to detect questionable bids: (1) there are two auctions of identical items at about the same time, where auction A ends before auction B; (2) the questionable bid is placed in auction B, even though an equal or lower bid in A would have been the highest bid in A; (3) the bidder who submitted the questionable bid in B, did not bid in A. Clearly, a buyer who places such a questionable bid could have done better by submitting a lower or equal bid to auction A, which terminates earlier than B.

Since questionable bids might appear for reasons other than shill bidding, Kauffman and Wood introduce additional requirements for identifying shill behavior. The criteria for detecting shill bids consist of: (1) shill bids are questionable bids; (2) shill bids are submitted by buyers who concentrate their bidding on the auctions of very few unique sellers; (3) shill bids are submitted earlier than average; (4) shill bids increase the current bid by more than the average bid; (5) shill bids are less likely to win the auction than the average bid. Of the more than 10,000 auctions examined by Kauffman and Wood (2003), 622 auctions (i.e. 6%) met their criteria for shill bidding. The authors also find that the probability of observing shill bidding in an auction increases when the minimum bid is relatively low, when the book value of the item is higher, when the auction duration is relatively long, and when the seller's other auctions appear to feature shill bidding. In a second paper, Kauffman and Wood (2005) show that when they identify an auction as having a shill bid, that auction tends to yield a higher price.

In another empirical study of shill bidding, Engelberg and Williams (2005) concentrate on finding evidence for the discover-and-stop technique described above. They examine almost 40,000 event-ticket auctions on eBay in September 2004, collecting additional information about the other auctions participated in by the buyers and sellers in that sample. In their 40,000 auctions, they find that 3% of all bids are discover-and-stop bids[23], and estimate that half of these, or 1.5%, are intentional shill bids. They also find that when bidders bid frequently on the same seller's auctions, the observed bids are more likely to be discover-and-stop shill bids.[24]

Hoppe and Sadrieh (2006) take a completely different approach than those of the empirical studies. They conduct a field experiment in which they auction blockbuster-movie DVDs and collectible-coin bundles. They auction three copies of each item simultaneously in three different treatments: (1) an auction with the minimum bid set to the lowest possible value, (2) an auction with the minimum bid set to about 60% of the book value, and (3) an auction with the minimum bid set to the lowest possible value, but with a later shill bid at about 60% of the book value. Both the literature on auction fever and the notion of avoiding the auctioneer's reserve-price fee suggest that the last treatment with the shill bid should result in the highest seller profit. In fact, the experiment shows that sellers in online auctions may have good reasons to use this latter setup. Although the observed prices are indistinguishable between treatments, seller profits are significantly higher in the two treatments with low minimum bids, because those strategies avoid paying additional fees for setting a high public reserve price. Interestingly, the effects are very similar no matter whether the auctioned item is in a "thick" market (blockbuster DVD) or in a "thin" market (unique collector's bundle).

---

[23] Engelberg and Williams define a discover-and-stop bid to be a bidder who bids twice, incrementally, within 10 minutes, and stops bidding as the second-highest bidder with evidence that the high bidder's bid is less than one increment higher.

[24] Shah *et al.* (2002) use data-mining techniques, on 12,000 eBay auctions of video-game consoles, to look for suspicious relationships between buyers and sellers, finding evidence suggestive that sellers often use multiple pseudonyms or confederates to execute their shilling. They point out the possibility of a shilling strategy of "unmasking," similar to Engelberg and Williams' discover-and-stop strategy, and propose (but do not execute) an analysis similar to that later performed by Engelberg and Williams.

Hoppe and Sadrieh (2006) observe no additional bid shading in the auctions with a shill bid; bidders in this environment would have no idea which auction (if any) involved shill bidding. By contrast, Kosmopoulou and De Silva (2005) verify the theoretical prediction of Chakraborty and Kosmopoulou (2004), that providing the subjects with knowledge about confederate bidding would induce them to increase their bid shading. In their laboratory experiment, subjects first participated in a number of ascending-price auctions without bid shilling, before the sellers were allowed to participate in bidding. Once the ability of the seller to participate was announced, the bidders' bid levels dropped, and average seller profit dropped from 97.5% to 88.9% of the item value. These findings highlight the seller's dilemma: sellers make more profit when there is no possibility of shilling, yet they always have an incentive to shill when bidders believe that shilling is not taking place. Once a seller's ability to shill is recognized by the bidders, bid levels drop, regardless of whether shill bids can actually be observed.

To summarize the results on shilling, we note that experimental research has shown that shill bids are effective tools for sellers in ascending-bid auctions, and that shill bids allow sellers to increase profits merely by escaping reserve-price fees. However, theory and experiments also suggest that sellers might be better off if the auctioneer could find some way to guarantee the absence of shill bidding. Empirical analysis suggests that shilling probably takes place in more than 1% but less than 10% of eBay auctions. EBay could probably reduce the rate of shilling not only by tracking close relationships between certain buyer and seller usernames, but also by changing its bidding system to thwart the discover-and-stop shilling strategy. For example, one possible remedy would be to eliminate the kinds of odd-number bids most effective in the discover-and-stop strategy, requiring all bids to be an even multiple of the current bid increment. Enforcement of rules against shilling can be quite tricky, especially as sellers become more sophisticated: a new service called XXLSell.com now provides a service (at least to German-speaking online-auction sellers) that automates the process of shilling, apparently using thousands of other XXLSell members' usernames to execute the shills in a way that is more difficult to detect.[25]

[25]Thanks to Henning Krieg for pointing out this interesting new development to us.

Section 4: Late and incremental bidding

Many researchers found that bids on eBay, where auctions run typically for a week, often arrive very near to the closing time – a practice called "sniping". For instance, in the sample of computer and antiques auctions with at least two bidders, Roth and Ockenfels (2002) found that about 50 percent of all auctions still have bids in the last five minutes, 37 percent in the last one minute, and still 12 percent in the last 10 seconds. Bajari and Hortacsu (2003) found that 32 percent of the bids in their sample are submitted after 97 percent of the auction has passed. Anwar et al. (2004) noted that more than 40 percent of the bids in their eBay sample are submitted during the final 10 percent of the remaining auction time. Simonsohn (2005) reported that in his sample almost 20 percent of all winning bids are placed with just one minute left in the auction, and Hayne et al. (2002) reported that bidding in the last minute occurs on average in 25 percent of their sample of 16,000 auctions. Regarding the whole distribution of the timing of bids, Roth and Ockenfels (2000) and Namazi (2005) observed that bid-submission times on eBay follow a power-law distribution with most bids concentrated at the closing time. Shmueli et al. (2005) added that the start of an auction also sees an unusual amount of bidding activity.[26]

At first glance, last-minute bidding of this sort cannot easily be reconciled with economic theory. As explained in section 2, eBay makes available a software bidding agent, called "proxy bidding," to make bidding simple for bidders without having to be constantly vigilant or online at the close of the auction. As a consequence, not the last bid (as in ascending-price auctions) but the highest bid wins, regardless of submission time. Furthermore, there is a risk involved in late bidding in online auctions. Because the time it takes to place a bid may vary considerably due to erratic internet traffic or connection times, last-minute bids have a positive probability of coming in too late (after the close of the auction).[27] eBay explains the simple economics of second-price auctions and the risks involved in late-bidding and comes to the conclusion: "*eBay always*

[26] There appear to be differences with respect to sniping frequencies across countries. Wu (2004) noted that there is much less sniping on eBay's Chinese platform Eachnet. However, one might speculate that Eachnet was not as mature as other platforms at the time of the study; the feedback score of most bidders was zero, and there were almost no competing auctions. Hayne et al. (2003a) reported that in their sample bidding occurs in the last minute of an auction with, for instance, 12 percent probability in United Kingdom and 36.5 percent probability in Sweden.

[27] In a survey of 73 bidders who successfully bid at least once in the last minute of an eBay auction, 63 replied that it happened at least once to them that they started to make a bid, but the auction was closed before the bid was received (Roth and Ockenfels 2002). Human and artificial bidders do not differ in this respect. The last-minute bidding service esnipe.com, which offers to automatically place a predetermined bid a few seconds before the end of an eBay auction, acknowledged that it cannot make sure that all bids are actually received on time by eBay.

recommends bidding the absolute maximum that one is willing to pay for an item early in the auction. (...) If someone does outbid you toward the last minutes of an auction, it may feel unfair, but if you had bid your maximum amount up front and let the Proxy Bidding system work for you, the outcome would not be based on time."

However, Ockenfels and Roth (forthcoming) demonstrated within an auction theoretic model that sniping on eBay could be a best response to a variety of strategies. In particular, inexperienced, "naïve" bidders might mistake eBay's proxy system for an ascending-price auction, and thus continually raise their bids to maintain their status as the high bidder. In an eBay style auction that closes at a predetermined deadline ('hard close'), bidding very late might be a best response to 'incremental bidding' (or multiple bidding) of this sort. That is, bidding very near the end of the auction would not give the incremental bidder sufficient time to respond, so a sniper competing with an incremental bidder might win the auction at the incremental bidder's initial, low bid. In contrast, bidding one's value early in the auction, when an incremental bidder is present, would win the auction only if one's value were higher than the incremental bidder's, and in that case one would have to pay the incremental-bidder's value.

Late bidding may also be a best response to other incremental bidding strategies for private-value model environments. One of these strategies is shill bidding by confederates of the seller in order to push up the price beyond the second-highest maximum bid. Barbaro and Bracht (2005), among others, argue that bidding late may protect a bidder from certain shill bidding strategies. Engelberg and Williams (2005) demonstrates how shill bidders may use incremental bids and eBay's proxy-bid system to make bidders pay their full valuations.

An additional reason for rational late bidding is given by Rasmusen (2003), where multiple bidding is caused by uncertainty over one's own private valuation (see also Hossain 2006). He argues within a game-theoretic model that bidders are ignorant of their private values. Thus, rational bidders may refrain from incurring the cost of thinking hard about their values until the current price is high enough that such thinking becomes necessary. Note that this, in turn, creates incentives for bidding late, because it prevents those incremental bidders from having time to acquire more precise information on their valuation of the object being auctioned.

Another rational reason for incremental bidding is that bidders may be reluctant to report their values, fearing that the information they reveal will later be used against them (see Rothkopf, Teisberg and Kahn 1990). While the highest maximum bid is kept secret on eBay, it sometimes happens that the winner defaults and that then the seller contacts the bidder who submitted the second-highest bid. If this bidder revealed his value during the auction, the seller can make a take-it-or-leave-it offer squeezing the whole surplus from trade. By bidding incrementally, private information can be protected – but only at the risk that a sniper will win at a price below one's value. Other papers refer to emotional factors as explanations for incremental bidding, such as 'auction fever' (Heyman et al. 2004), escalation of commitment and competitive arousal (Ku et al. 2005). Another explanation along these lines is the 'pseudo-endowment effect' (Wolf, Arkes and Muhanna 2005), which posits that temporarily being the high bidder during an auction increases the bidder's value. Note that not only is late bidding a good strategy to avoid incremental bidding wars with *other* emotional bidders, but also that late bidding may also serve as a *self*-commitment strategy to avoid how one's own bids are affected by auction fever and endowment effects.

The evidence in the laboratory and the field indicates that incremental bidding is common, and that sniping is likely to arise in part as a response to incremental bidding. Wilcox (2000) indicates that the average bidder submits 1.5 to 2 bids. Ockenfels and Roth (forthcoming) report that 38 percent of the bidders submit a bid at least twice. Among these bidders, the large majority submits a new bid after being outbid. In particular, 53 percent of the last bids of incremental bidders are placed after the previous bid was automatically outbid by eBay's proxy bidding agent (i.e., by another bidder's proxy that was submitted earlier in the auction), 34 percent are placed after the previous bid was outbid by a newly submitted proxy bid of another (human or artificial) bidder, and only 13 percent are placed by the current high bidder (so that the current price is not changed). Bids per bidder increases with the number of other bidders who bid multiple times in an auction, which suggests that incremental bidding may induce bidding wars with like-minded incremental bidders.[28] In a regression study using eBay field data, Wintr

[28] They also note that naive English-auction bidders may also have an incentive to come back to the auction close to the deadline in order to check whether they are outbid. However, the data indicate that among those bidders who submit a bid in the last ten minutes of an eBay auction, one-bid bidders submit their bid significantly later than incremental bidders. The data also reveal that bidders with a larger feedback score tend to submit less bids per auction, suggesting that incremental bidding is reduced with experience. However, in a study by Hayne et al. (2003) the bidders who submitted multiple bids had a higher average feedback score than the average for all bidders.

(2004) found that the presence of incremental bidders leads to substantially later bids, supporting the view that sniping is reinforced by incremental bidding.

Ariely et al. (forthcoming) investigated the timing of bids in a pure private-value laboratory setting. They observed that early bids are mostly made in incremental bidding wars, when the low bidder raises his bid in an apparent attempt to gain the high bidder status, while late bids are made almost equally often by the current high bidder and the current low bidder. That is, late bids appear to be planned by bidders regardless of their status at the end the auction.

Incremental bidding is not the only possible cause for late bidding. Roth and Ockenfels (2002) and Ockenfels and Roth (forthcoming) demonstrate that there can be equilibria where all bidders submit only one bid late in the auction, even in purely private-value auctions and even though this risks failing to bid at all. This kind of equilibrium can be interpreted as collusion against the seller, because it has the effect of probabilistically suppressing some bids, and hence giving higher profits to the successful bidders. Several researchers tested the implications of the model, but the model could generally not be supported. Using eBay field data, Bajari and Hortacsu (2003) could not statistically confirm whether early bids lead to higher final prices. Hasker et al. (2003) as well as Wintr (2004) could not find evidence that the distribution of final prices is different for winning snipes and winning early bids on eBay. In a controlled field experiment, Gray and Reiley (2004) found 2.54 percent lower prices when the experimenter submitted the bid just 10 seconds before the end of the auction compared to when the bid was submitted several days before the end, though the difference was not statistically significant. Finally, in their laboratory study, Ariely et al. (forthcoming) found that when the risk of sniping is removed, the amount of late bidding goes up. This evidence also contradicts the ‘implicit collusion’ explanation, for late bidding should decrease when there is no chance of suppressing bids through sniping. However, most of the studies reported substantial amounts of multiple bidding. This again suggests that parts of the sniping behavior can be interpreted as a response to incremental bidders.

Another explanation for incremental bidding without positing inexperience on the part of the bidders is to note that, if an auction is common-value rather than private-value, bidders receive information from others’ bids that causes them to revise their willingness to pay. In general, late bids motivated by information about common values arise either so that bidders can incorporate into their bids the information they have gathered from the earlier bids of others, or so bidders can avoid giving information to others through their own early bids. Bajari and Hortacsu (2003) formalize this idea in a symmetric common value model; Ockenfels and Roth (forthcoming) give an example of equilibrium sniping in a simple common-value model with asymmetrically informed bidders.

Roth and Ockenfels (2002) provides survey evidence, and Ockenfels and Roth (forthcoming) provides field evidence which supports the common value explanation. They show that there is less last-minute bidding on eBay computer auctions than on eBay antiques auctions, which supposedly possess more common value elements. However, the fact that Ariely et al. (forthcoming) observed substantial sniping in the laboratory for a pure private-value context strongly suggests the common-value explanation that bids are interpreted as value signals does not entirely explain the motivations for sniping behavior.

Another direction for explaining late and multiple bidding is based on the multiplicity of listings of identical objects, which may create incentives to wait until the end of an auction in order to see how prices develop across auctions. Peters and Severinov (forthcoming) propose a model with simultaneously competing auctions and argue that late bidding is consistent with this model. Stryszowska (2005a; see also 2005b) models online auctions as dynamic, private-value, multi-unit auctions. By submitting multiple bids, bidders coordinate between auctions and thus avoid bidding wars. In one class of Bayesian equilibria, multiple bidding also results in late bidding, even when late bids are accepted with probability smaller than one. Wang (2003) shows theoretically that in a twice–repeated eBay auction model, last-minute bidding is in equilibrium and offers some field evidence for this. Anwar et al. (2004) provide evidence suggesting that eBay bidders tend to bid across competing auctions and bid on the auction with the lowest standing bid. This seems to support the idea that the incentives to bid late are amplified when there are multiple listings of the same item. However, it is somewhat surprising how little empirical work on competing online auctions has been conducted.[29]

Some observers of eBay believe that the amount of sniping will decrease over time because it is mainly due to inexperience and unfamiliarity with eBay's proxy bidding system. This

[29] Two other studies (Zeithammer 2004, and Arora et al. 2003) theoretically and empirically study bidding behavior in sequential online auctions but they do not address the issue of bid timing within a given auction. Vadovic (2005) studies dynamic auctions in which bidders coordinate who searches for outside prices and shows that bidders with low search costs tend to bid late.

section showed, however, that there are a variety of rational, strategic reasons for sniping. It is a best response to naïve and other incremental bidding strategies, and can even arise at equilibrium in both private-value and common-value auctions. In fact, Wilcox (2000), Roth and Ockenfels (2002), Wintr (2004), and Ariely et al. (forthcoming) observed, both in laboratory and field studies, that more experienced bidders snipe more often than less experienced bidders.[30] Thus, as long as the auction rules remain unchanged, it seems likely that late bidding will remain a persistent phenomenon on eBay.

Section 5: The buy-now option

A feature of online auctions increasing its share of market transactions is the *buy-now option*. This option, also known as *buyout option* allows any bidder to end the auction early at a *buy-now price* previously specified by the seller. When a bidder exercises a seller's buy-now option, he pays the specified price (some authors prefer to call this price the *buy-out price* or the *buy price*), in order to receive the item immediately and shut out other bidders. On eBay, the buy-now price is called the *Buy it Now price*, whereas on Yahoo it is called the *Buy price*. The buy-now option used by eBay and Yahoo differ by more than just the name. The Buy-it-Now option on eBay is a *temporary buy-now option*, available only so long as no bid has yet been placed on the item.[31] The Buy-it-Now option disappears after the first bid. In contrast, the Buy-price option on Yahoo is a *permanent buy-now option* and available for the entire duration of the auction. Other online auctioneers use these options as well: a survey by Matthews (2005) shows that temporary buy-now options are available on eBay, LabX, and Mackley & Company, whereas Yahoo, uBid, Bid or Buy, MSN, and Amazon offer permanent buy-now options.[32] The practice appears to be much more common in the online world than in traditional English auctions. We are not aware of any documented use of buy-now prices in traditional auctions, perhaps because live auctions usually take place on much smaller timescales (seconds versus days).

---

[30] Simonsohn (2005) investigated the consequences of such lateness on the strategic behavior of sellers. The idea is that because many bidders snipe, an auction's ending-time is likely to influence the number of bidders it receives. In fact, he found that a disproportionate fraction of sellers set the ending-time of their auctions to hours of peak-demand.

[31] If the seller has set a secret reserve price, the Buy it Now option remains active until the reserve price is overbid.

[32] This practice developed even before commercial auction websites. Lucking-Reiley (2000a) documents the use of "buyout prices" in English auctions run by individuals on online newsgroups before the advent of eBay.

The buy-now option has become increasingly popular with sellers. According to Matthews and Katzman (2006), of the items for sale on eBay, 30% were listed with a buy-now option in the first quarter of 2001 and 35% in the second quarter of 2001. By December 2001, 45% of items for sale on eBay were listed with a buy-now option. Hof (2001) generally agrees with these findings, reporting that about 40% of the items for sale on eBay had the buy-now option towards the end of 2001. Kane (2002) reports that in the second quarter of 2002, 33% of eBay items listed worldwide had the buy-now option and accounted for 19% of gross merchandise sales. The economic relevance of the buy-now option can easily be seen in eBay's quarterly reports. The reports for 2005 show that eBay's fixed-price sales, consisting largely of Buy-it-Now purchases, accounted for $3.2 billion (30% of the gross merchandise sales) in quarter 1, $3.2 billion (29%) in quarter 2, $3.4 billion (32%) in quarter 3, and $4.0 billion (34%) in quarter 4.

The existence and increasing popularity of the buy-now option are puzzling from the point of view of auction theory. An auction's primary benefit is that it relieves the seller of the job of determining an item's price, instead allowing bidders to determine the price by competing with each other. Introducing a buy price could potentially decrease a seller's revenue, because when exercised it rules out the possibility of higher prices reached by competitive bidding. If exercised by a bidder with less than the highest value, the buy price can similarly reduce efficiency. Given the potential inefficiency and loss of revenue, why is the buy-now option so popular with sellers?

5.1 Explaining the buy-now option with risk-aversion

One of the first proposed explanations for the observed popularity of the buy-now option is the risk aversion of bidders or sellers. Budish and Takeyama (2001) show that adding a permanent buy-now option to an ascending auction can increase the seller's revenue, in a model with two risk-averse bidders with only two possible valuations. Reynolds and Wooders (2003) extend the result to a continuous uniform distribution of valuations, demonstrating that it holds for both types of buy-now options. They show that optimally chosen buy-now prices are never exercised in equilibrium when the bidders are risk-neutral. In contrast, when the bidders are risk-averse, the optimal buy-now prices are exercised with a positive probability, providing insurance value to risk-averse bidders and increasing the risk-neutral seller's revenue.

Though a buy price may in principle allow a lower-value bidder to shut out the highest-value bidder, Hidvégi, Wang and Whinston (2006) show that in ascending auctions with a permanent buy-now option and uniformly risk-averse bidders, such displacement of the high valuation bidder will not occur in equilibrium. Intuitively, with a permanent buy-now price set optimally by the seller, no bidder immediately jumps to the buy-now price. Instead, bidders with valuations that are high enough to accept the buy-now price will first engage in straightforward bidding until the current bid level has reached a certain threshold. If bidding reaches some bidder's threshold, that bidder ends the auction by accepting the buy-now price. Assuming all bidders have exactly the same degree of risk aversion, their thresholds decrease monotonically in their valuations (i.e. the bidder with the highest valuation will be the first to jump to the buy-now price). This monotonicity of threshold values ensures an efficient outcome of the auction, and the optimally chosen buy price yields at least as much expected revenue as an auction without a buy price.[33]

Hidvégi, Wang and Whinston (2004) note that their efficiency and neutrality results break down if the permanent buy-now option is replaced by a temporary buy-now option. Temporary buyout options do not allow for the type of threshold strategies discussed above. Both efficiency and seller revenue are lower than in an auction with a permanent buy price. Permanent buy prices also produce higher revenues than temporary buy prices in the risk-averse-bidder model of Reynolds and Wooders (2003) and the impatient-bidder model of Gupta and Gallien (2005).

Using a slightly different approach and focusing on a temporary buy-now option, Mathews and Katzman (2006) show that buy-now prices may increase expected utility for risk-averse sellers facing risk-neutral bidders. The intuition here is simply that a risk-averse seller may be willing to give-up part of the expected auction revenue to reduce the volatility of auction revenue. In an extreme case, an infinitely risk-averse seller can choose a buy-now price low enough for even the lowest-valued buyer type to accept. This guarantees immediate trade at a fixed price. A seller with less extreme risk aversion will choose a higher buy-now price, earning higher expected revenues with nonzero variance. The buy price produces higher expected utility for the seller even though it may result in an inefficient allocation.

[33] Interestingly, the risk-averse bidders in this model do not receive higher expected utility from the presence of a buy-now price. Though they receive some benefit from reduced uncertainty, the seller manages to extract the added bidder surplus with an optimally chosen buy-now price.

5.2 Explaining the buy-now option with impatience and other transaction costs

An alternative explanation for the prevalence of the buy-now option is the impatience of the trading agents. An impatient bidder may be willing to pay a premium to receive the item quickly. Similarly, an impatient seller may be willing to accept a lower price to end the auction early. Indeed, eBay cited impatience as a reason for introducing the Buy-it-Now price (Dennehy, 2000). Mathews (2003) studies both bidder and seller impatience in an independent-private-valuations auction with a temporary buy-now option. He shows that impatience on either side of the market creates incentives for the seller to set a buy-now price that would be exercised with positive probability. The optimal buy-now price increases directly with bidders' impatience, inversely with seller's impatience, and inversely with the number of bidders.

Gupta and Gallien (2005) also examine the case of "time sensitive" (i.e., impatient) bidders in an independent-private-value auction with bidders arriving at the auction via a Poisson process. In addition to comparing temporary and permanent buy-now options, they investigate the theoretical possibility of a "dynamic" buy-now price whose level can be changed by the seller during the course of the auction. The authors first solve for the Nash equilibrium in the bidders' strategies for each of the auction formats, then use numerical simulations to derive results on the optimal seller choice. The results show that a dynamic buy-now price barely improves seller utility compared to a static one, which perhaps explains the lack of dynamic buy prices observed in the field. The simulations also show that a permanent buy-now option can enhance revenues far more than temporary buy-now option can.[34]

In the models discussed so far, impatience makes auction participation costly for a bidder relative to a fixed-price purchase. Wang, Montgomery, and Srinivasan (2004) suggest that a consumer may have other substantial transaction costs associated with participation in an auction. Amongst others, they describe the cognitive effort that is necessary to observe, plan, and execute the bidding in an auction.[35] In the presence of these bidder transaction costs, buy prices produce benefits similar to those derived in the above models of bidder impatience.

[34] Interestingly, the model also predicts increased late bidding when a permanent buy price exists. This result is related to the fact that the auction price approaches the permanent buy price towards the end of the auction.

[35] See Engelbrecht-Wiggans (1987) for additional discussion of auction entry costs.

5.3 Explaining the buy-now option with a sequence of transaction opportunities

An alternative explanation for the observed frequent use of buy-now prices in online auctions rests on the idea that sellers may use them to optimize revenues intertemporally when identical units of the item come up for sale at different points in time. Kirkegaard and Overgaard (2003), for example, examine a sequence of single unit auctions, run either by the same seller or by multiple sellers. In the latter case, the first seller can increase her revenues and decrease the revenues of the subsequent sellers by choosing the optimal temporary buy-now price. In the former case, when the seller is a monopolist, the optimal strategy prescribes not to use the buy-now option for the first auction, but to announce its use for later auctions. Thus, the buy-now price in this model is a valuable instrument to the seller even though all agents are risk-neutral and no one is impatient.

Etzion, Pinker, and Seidmann (2003) describe a similar result for the case of a multi-unit seller who faces a stream of randomly arriving bidders and can sell any number of items either in unrestricted auctions or with a permanent buy-now price. The seller uses both mechanisms to effectively price discriminate between buyers with low versus high willingness-to-pay. The former would earn negative surplus at the buy-now price, so they merely participate in the auction. The latter, who could earn positive surplus at the buy-now price, generally do not do so immediately: they bid in the auction at first, and only later consider accepting the buy-now price. In a closely related model, Caldentey and Vulcano (2004) derive a similar equilibrium with a more complex model of bidders' utility.

Many online auctions represent clearance sales of overstock items. That is, both on eBay and on merchandise sites such as uBid.com, auctions take place for items that failed to sell at a posted price in ordinary sales channels. Obviously, if buyers rationally expect the possibility of purchasing an item in a later overstock auction, that might affect their initial purchase decision. Note that this possibility can be modeled as a multiple-item auction with a temporary buy-now price. If all items sell immediately at the buy-now price, no auction takes place, but any "leftover" items are sold at auction.

Bose and Daripa (2006) analyze this scenario. They model buyers' valuations rather unconventionally: A buyer either has a pre-specified high valuation or a valuation randomly drawn from a continuous distribution strictly lower than the high valuation. The seller's problem is to price discriminate by making the buy-now option attractive enough to the high-value buyers so that they do not wish to wait for the auction. Bose and Daripa show that the seller cannot perfectly price discriminate; the optimal mechanism involves a temporary buy-now price set low enough to be attractive to some of the low-valuation buyers as well.

5.4 Empirical and experimental evidence

The growing body of alternative theoretical explanations calls for empirical and experimental testing. Some preliminary research is available, but much work remains to be done, both on drawing out the testable implications that might distinguish the competing theories, and on collecting empirical data to test them.

The earliest empirical studies of buy-now prices have mainly generated descriptive data about the practice. Mathews (2003), for example, reports the prevalence and execution of the buy-now option in two specific product categories on eBay (software for the Sony PlayStation). He finds the buy-now option available for more than half of the items (59%). The option was exercised in about 27% of the cases in which it was available. In about 62% of these cases, the buy-now price was below the auction price.

Reynolds and Wooders (2003) also provide frequencies with which the buy-now option has been used on eBay and Yahoo. They sample a total of 31,142 eBay auctions and 1,282 Yahoo auctions in the categories of automobiles, clothes, DVD-players, VCRs, digital movie cameras, and TV-sets. 40% of eBay auctions used the temporary Buy-it-Now option, while 65% of Yahoo, auctions used the permanent Buy-Price option. Since the permanent buy-now option of Yahoo is chosen more frequently than the temporary buy-now option of eBay, this gives some empirical support to those theories predicting that the permanent buy-now option may be more beneficial for sellers than the temporary option.

Durham, Roelofs, and Standifird (2004) examine a sample of 138 auctions of American silver dollars on eBay. They find that the 41 auctions using a buy-now price (temporary since it is on eBay) result in significantly higher selling prices (on average $10.27) than the auctions without a buy-now price (on average $9.26). From these findings, the authors conclude that the buy-now option increases seller revenue, consistent with the idea that there are some risk-averse or impatient bidders willing to pay a premium to guarantee a win or end the auction early.

Anderson, Friedman, Milam, and Singh (2004) collected data on about 1000 Palm Pilot Vx auctions on eBay. The goal of the study is to identify and relate typical seller profiles to typical seller strategies. One of the most interesting findings is that the high-volume sellers typically use a buy-now price in combination with a very low minimum bid.

Hendricks, Onur, and Wiseman (2005) analyze data from almost 3,000 Texas Instruments TI-83 Graphing Calculator auctions on eBay, where 831 (roughly 30 percent) offered a buy-now price. The auctions with buy-now prices appeared to produce significantly higher revenue than the auctions without buy-now prices. However, by contrast with the behavior of the high-volume sellers studied by Anderson et al., a great majority of the auctions with a buy-now price also featured a high minimum bid, 90% or more of the buy-now price. Because the buy-now auctions also tended to have higher minimum bids than the non-buy-now auctions, it is hard to tell whether the cause was the buy-now price or the higher minimum bid (see Section 3)..

A number of studies have used experimental methods to study the buy-now option. Standifird, Roelofs, and Durham (2005) report a field experiment in which they auctioned 84 American Eagle silver dollars on eBay to study the impact of varying buy-now prices. Surprisingly, they find that eBay buyers hardly made use of the buy-now option, even when the buy-now price was substantially below the prevailing market price. The authors suggest that buyers may be reluctant to use the buy-now option, in order not to forfeit the entertainment benefit associated with the participation in an eBay auction.

Shahriar and Wooders (2005) report laboratory experiments examining whether sellers can profitably use the buy-now option in independent-private-value auctions with risk-averse bidders. Common value auctions, in which the buy-now option theoretically creates no advantage, are also examined in a control treatment. The authors find that suitably chosen buy-now prices raise sellers' revenues in both treatments. They speculate that the unexpected positive effect of the buy-now price on the sellers' revenues in the common-value treatment may be due to winner's-curse-type overbidding.

Using a similar experimental design, Seifert (2006) reports a strong interaction effect between the market size and the revenue impact of a temporary buy-now option. The experiment shows that buy-now prices have the expected positive impact on sellers' revenues with five active bidders, but the effect is lost when the number of active bidders falls to three.

All in all, the buy-now option remains rather poorly understood. We see few robust findings, and little ability to discriminate between the different proposed theoretical explanations.

As a relatively new economic phenomenon, the buy-now option clearly remains an area in with many opportunities for exciting empirical research.

Section 6: Parallel markets and other outside options

Before the advent of electronic markets, most auctions took place as fairly isolated events. The items sold in an auction were often not easily available elsewhere, at least not at feasible time and travel costs. Furthermore, bidders were constrained to being at one auction at a time and usually could not bid on several auctions in parallel.[36] In general, the only alternative for a bidder, who did not receive the item in an auction, was to wait for a subsequent auction. Hence, early auction literature has been fairly extensive on sequential auctions[37], but has almost completely ignored settings in which bidders can simultaneously participate in multiple auctions or use other sales channels to purchase the item. In such settings, the seller in an auction is no longer a monopolist, but faces competition by the sellers in other auctions or markets. The bidders do not only have an entry choice, as in isolated auctions with endogenous entry decisions, but have an outside option in the alternative sales channel.

To see the empirical relevance of online auctions with parallel markets, it suffices to start up any Internet browser and open one window on a typical online auction (e.g. eBay or Yahoo!) and another window on a typical shopbot (e.g. mysimon.com or shopping.com). For hundreds of products, there are numerous auctions running in parallel and numerous online shops offering fixed price sales at the same time. Auction sellers are often in a highly competitive environment that is substantially different from the monopoly position they are assumed to have in classical auction theory. Buyers must choose the type of market (e.g. auction or posted offer), the specific instance of that type (which auction or which posted offer), and – if necessary – the bid to submit. These are all more complicated strategic decision situations than those that have typically been studied in the auction theory literature.

The fact that the strategic interaction quickly turns overly complicated, when competition and outside options are added to the standard auction models, has limited the amount of

[36] Telephone bidding clearly presented the first step towards bidder presence at multiple auctions. However, since telephone bidding is quite costly and only a small portion of all bidders use it, the auctioneers often plan their auctions in a way to reduce parallel auctions of the same category of goods.

[37] See the overviews in Klemperer (1999) and Krishna (2002).

theoretical work on the topic. McAfee (1993) shows that equilibria in games in which sellers compete by offering different direct mechanisms, may not be feasible in general, due to the possible non-convexity and discontinuity of the sellers' profit functions. However, under some restrictive assumptions, McAfee (1993) shows the existence of an equilibrium in which all sellers offer second-price auctions with the reserve prices set to their marginal cost.

Peters and Severinov (1997) use a limit equilibrium concept to explicitly characterize the symmetric equilibrium that arises when the seller's competition is restricted to choosing only amongst auctions (instead of choosing any arbitrary direct mechanism). The author's approach also allows them to examine the efficiency aspects of the auction competition equilibria. They show that the efficiency of the auction competition market depends on the seller's ability to advertise his reserve prices and on the timing of the buyers' knowledge of the own valuation. An efficient market performance, in the analyzed setting, is achieved when sellers can advertise their reserve prices and buyers learn about the realization of their valuation only after having chosen a specific auction. If buyers are informed of their valuations before they select an auction, the reserve prices are driven down to the sellers' marginal costs and there is inefficiently excessive entry. All these results, however, are derived under the restrictive symmetry assumption that buyers randomize their purchases over all sellers who offer the same conditions. This is a restrictive assumption; it excludes all possible equilibria in which buyers systematically sort amongst sellers.

Building on the two studies discussed above, Peters and Severinov (forthcoming) characterize a dynamic adjustment mechanism that provides a perfect Bayesian equilibrium for a market where sellers compete by offering different reserve prices in their independent ascending second-price auctions. Buyers bid in multiple rounds, costlessly and independently adjusting their bids and moving from one auction to the other, whenever their previous bid was not successful. The perfect Bayesian equilibrium derived for this decentralized trading institution induces an efficient set of trades at a uniform trading price.

The strength of this equilibrium is that the decentralized equilibrium bidding has extremely low informational requirements and neither depends on a buyer's beliefs about the other buyers' valuations, nor on the number of buyers and sellers. The equilibrium bidding rule only requires that any buyer, who currently does not hold a high bid, should bid in the auction with the lowest current high bid, raising the high bid as slowly as possible and exiting the market if the calculated next bid is greater than one's valuation. This *minimal increment bidding* is a feasible strategy for most of the existing online auctions, because the only two pieces of information that a bidder needs to follow this strategy are generally available: (1) the information whether one's bid is currently the high bid of an auction and (2) knowledge of which high bids are currently associated with each of the auctions. Note, however, that while minimal increment bidding is feasible, it is hardly ever observed in online auctions.

Given that bidders use minimal increment bidding, Peters and Severinov (forthcoming) show that in equilibrium all sellers will set their reserve prices equal to their marginal costs if the number of traders in the market is sufficiently large. Under these circumstances, the induced equilibrium of the parallel auctions market is efficient and sequentially optimal at every stage.[38]

The empirical evidence concerning parallel online auctions is mixed. Tung, Gopal, and Whinston (2003) tracked simultaneous online auctions of identical consumer electronic items, controlling for seller reputation and quality. They report large price disparities and, hence, substantial arbitrage opportunities across the auctions. They observed only a few *cross-bidders*, i.e. bidders switching from one auction to another. Interestingly, they note that none of the cross-bidders they identified ever succeeded to buy an item. Furthermore, they observed that bidders who were outbid in one auction did not switch to another, even though their unsuccessful bid was substantially higher than the winning bid of the latter auction. All this seems to indicate that bidding in online parallel auctions little resembles the minimal-increment bidding required in the equilibrium specified by Peters and Severinov (forthcoming).

There is, however, also empirical evidence in support of the Peters and Severinov (forthcoming) equilibrium analysis. Anwar, McMillan, and Zheng (2006) collected data from competing eBay auctions for CPUs. Controlling for all auction parameters, three samples were generated, which differed only in closeness of the auctions ending times: same day, same hour, and same minute. As suggested by minimal increment bidding, bidders tend to bid on the auction with the lowest high bid. Furthermore, the smaller the difference between the ending times, the more cross bidding is observed. Finally, the authors show that bidders using cross-bidding

[38] Burguet and Sákovics (1999) point out that efficiency breaks down in a setting in which the number of competing sellers is small enough to allow for strategic interaction. Examining the case of two auction sellers, who strategically set their reserve prices in a mutual best-response, they show that equilibrium reserve prices are well above the marginal cost, inefficiently excluding some bidders.

strategies on average pay only 91% of the price paid by those not using cross bidding. This evidence indicates that the bidders understand the strategic situation of competing auctions and react to it in an appropriate way. However, the fact that bids are far from being increased at the smallest increment indicates that bidders do not actually employ the predicted minimal increment bidding. Additionally, the price differences among auctions of identical goods seem to suggest that the observed behavior in the competing auction markets does not perfectly match the equilibrium behavior predicted by Peters and Severinov (forthcoming).

Stryszowska (2005a) takes a slightly different approach to competing auctions. She analyzes two simultaneous, second-price, private-value auctions for an identical item.[39] The interesting twist in the derived equilibria is that bidders may use early, low, non-consequential bids to identify themselves and to coordinate across the auctions. This early coordination effort can explain the multiple bidding that is frequently observed in online auctions. Furthermore, Stryszowska (2005a) shows that in some equilibria of the game, all crucial bids arrive early, thus inducing an efficient allocation with identical prices in both auctions. However, in another type of equilibrium, in which bidders send last minute bids, the auction outcome is inefficient and prices may be dispersed.

A completely different path for modeling the outside option is taken by Reiss (2004). Instead of examining parallel auctions that mutually affect one another, Reiss (2004) aggregates whatever outside opportunity a bidder has into a simple payoff value that the bidder receives if he is not successful in the auction. This payoff, for example, can represent the utility gain of buying the item at a posted-offer shop. Given the model with a bidder outside option, Reiss (2004) shows that the optimal auction reserve price is decreases as the outside option increases. In a related model, Kirchkamp, Poen and Reiss (2004) implement the bidder outside option as fixed payments to all the bidders who are not successful in the auction. In the experiments, the theoretical finding that increasing the outside options decreases aggressive bidding is reproduced. But, the experiments also show that bidders in second-price auctions manage to fully expropriate their outside option, while those in the first-price auctions fail to do so. Hence, outside options

[39] Instead of examining two parallel auctions, some authors analyze the situation with a sequence of two auctions that is known to the bidder. Zeithammer (2003) shows that in such a setting, bidders will "bargain-hunt" and reduce their bids, if a more preferred item is next up for sale. A similar approach can be found in Reiss and Schoendube (2002) and Brosig and Reiss (forthcoming).

seem to amplify the well-documented effect that first-price auctions generate more revenues than second-price auctions. In terms of efficiency, however, neither auction type is strongly affected, leaving the differences between the two designs insubstantial.

Section 7: Multi-item auctions

With few exceptions, previous sections dealt primarily with single-item online auctions. However, almost all online auction platforms also offer multi-item auction mechanisms. In recent years, multi-item auctions have received increasing attention, mainly because of their accelerated use in B2B (business to business) commerce and government allocation procedures. Yet, both the theoretical and empirical literature is less developed and contains only few general results. This is partly because when items are heterogeneous or bidders demand multiple items, new difficulties such as market power and strategic and computational complexities arise. Here we present the standard auction mechanisms for selling multiple items, and we discuss some central intuitions as well as empirical evidence on bidding behavior. For an in-depth overview of multi-item auction theory, see Milgrom (2004).

7.1 Standard multi-unit auction mechanisms

For selling multiple units of one item (such as car tires, financial securities, energy products, environmental permits, etc.), there are, analogous to the single-object case, four standard mechanisms: the descending-price auction, the sealed-bid pay-as-bid auction, the ascending-price auction, and the sealed-bid uniform-price auction, plus some variations and extensions.[40]

In the *ascending-price* multi-unit auction, the price gradually increases while bidders indicate how many units they want at each price. The final price is set and the auction closes when aggregate demand equals the number of units supplied. All bidders pay the same final price, which is the price at which the auction closed. In the corresponding sealed-bid version, bidders independently submit a whole demand curve. That is, each bidder indicates how much he is willing to pay for the first unit he acquires, the second unit, etc. Then, the outcome of the auction is determined by finding the first price at which aggregate demand equals supply. All

[40] All auction mechanisms considered here are simultaneous auctions; for a brief discussion of sequential auctions see section 8.3.

bids above or equal to this price win, and bidders must pay their bid price.[41] In the *sealed-bid uniform-price* auction, all units have the same price: the market-clearing price.

Note that the ascending-price and the sealed-bid, uniform-price mechanisms enforce uniform prices for all units sold. Other auction formats endogenously promote 'similar prices for similar objects' by encouraging arbitrage. The best-known auction mechanism in this category is the *simultaneous ascending auction* (SAA) developed by Milgrom, Wilson, and McAfee, who proposed the design in the context of the US radio spectrum auction. All items, which may or may not be identical, are simultaneously offered in different auctions. Bidding on all objects takes place simultaneously in rounds subject to an activity rule (see section 8.2). Bidders observe prices throughout the auction, and this information allows them to arbitrage among substitute licenses, and to piece together complementary packages. The auction ends when a round passes with no new bids on any licenses (see, e.g., Cramton 2002 and 2004 for more details). For multiple, divisible items, Ausubel and Cramton (2004) advocate the *simultaneous clock auction*, where a price clock for each divisible good indicates its tentative price per unit quantity. Bidders express the quantities desired at the current prices, and the price is then repeatedly increased by an increment until demand is made equal to supply, at which point the tentative prices and assignments become final. This auction also yields similar prices for similar items by encouraging arbitrage. On the other hand, the next two multi-unit auction formats are discriminatory; identical units are sold at different prices.

In the *decreasing-price* multi-unit auction, the price gradually decreases while bidders indicate the price they are willing to buy one or more units. At each price, bidders are informed about the supply left at that point. The auction closes when no supply is left. Each winner pays the price at which he indicated he was willing to buy. In the *sealed-bid pay-as-bid* auction, bidders independently submit a demand curve. Every winning bidder pays his bid for each unit, provided that the bid is above the clearing price.

Another multi-unit auction with non-uniform prices was proposed in the seminal paper by Vickrey (1961). Suppose there are k units for sale. As before, the highest k bids are accepted, but the pricing rule of the *Vickrey-auction* determines that for the k^{th} unit awarded, bidders have

[41] When the number of units is an integer, the clearing price may be the lowest accepted bid or the highest rejected bid. We also note that if the units being sold are not substitutes in the eyes of the bidders, then market clearing prices can fail to exist; see Milgrom (2004).

to pay the amount of the *k*th highest losing bid.[42] This rule generalizes Vickrey's second-price auction rule for the single-item auction, where the winner pays the largest losing bid, to the multi-unit case. In fact, analogous to the single-object case, all bidders have a dominant strategy to bid true values for all units.

Search engines such as Google typically use multi-item auctions, distantly related to the Vickrey auction, to sell online ads. Search results are typically shown along with sponsored links, which in turn are shown in decreasing order of bids. If a user of the search engine then clicks on an ad in position k, that advertiser is charged by the search engine an amount equal to the next highest bid, i.e., the bid of an advertiser in position $k + 1$. Because there are multiple positions available, there are many winners, and each winner pays the next highest bidder's bid. Edelman et al. (2005) show that this 'generalized second-price' auction generally does not have an equilibrium in dominant strategies. But, it has a unique ex-post equilibrium, resulting in the same payoffs as the dominant strategy equilibrium of the Vickrey auction (see Varian 2006 for another analysis of Google's 'position auction').

7.2 Bid shading and demand reduction in multi-unit auctions

When bidders do not demand more than one unit, the analysis of the single-item case straightforwardly generalizes. For instance, bidshading will occur in sealed-bid pay-as-bid auctions (reflecting the trade-off between the probability of winning and the surplus from winning), and 'truth-revealing' in the ascending-price and the uniform-price auction. In fact, just as in the single-item case, the standard auction rules are revenue-equivalent under some appropriate assumptions. However, with multi-unit demand, bidding incentives can be quite different, and revenue equivalence fails to hold.

Maybe the most important intuition from the literature is that uniform-price auctions do not share the desirable properties of the second-price auction in the single-item case. The reason is that if a bidder can demand more than one unit, there is a positive probability that his bid on a second or later unit will be pivotal, thus determining the price for the first and possibly other units. With discrete goods, the bidder will bid his true value on the first unit, but strictly less on

[42] Suppose, e.g., that there are three bidders A, B, and C competing for three units of an object. Bidder A bids 14, 10 and 2, bidder B bids 12, 9 and 0, and bidder C bids 8, 5 and 4, respectively. Then bidder A is awarded two units and bidder B one unit. Bidder A pays 17 (= 9 + 8) for both units, and bidder B pays 9 for his unit.

all subsequent units. As a consequence, in equilibrium, bidders understate their values, or (equivalently) reduce demand quantities which hampers revenue and efficiency.[43] Furthermore, uniform-price auctions typically facilitate (tacit or explicit) collusion. Suppose the bidders agree on a collusive agreement and each bidder bids higher prices for smaller quantities than his collusively agreed share. Then, if any bidder attempts to obtain more, all bidders would have to pay high prices. This stabilizes collusion. So, a key concern with uniform-price auctions is the possibility of low price equilibria.

Several field studies provide direct evidence of strategic demand reduction and collusive behavior in electronic auction markets, such as in the German auction of GSM spectrum (Grimm et al. 2003), in the Austrian auction of third generation mobile wireless licenses (Klemperer 2004), in the FCC's Nationwide Narrowband Auction (Cramton 1995), in the UK electricity market (Wolfram 1998), and in the California electricity market (Borenstein et al. 2002). This field evidence is strongly supported by laboratory evidence (e.g., Kagel and Levin 2001, Engelmann and Grimm 2004) and controlled field experiments (List and Lucking-Reiley 2000). It has also been shown that, in line with theory, the amount of demand reduction decreases with the number of bidders (Engelbrecht-Wiggans et al. forthcoming). There is, however, little research on multi-unit bidding and demand reduction in online auctions (but see the chapter by Morgan et al. 2005 in this Handbook on competing auction research).

The two most common online multi-unit formats are the "Yankee auction" (as used by Onsale.com and also called "Multiple Item Progressive Electronic Auction"; see Bapna et al. 2000) and eBay's "Dutch auction" as used in the US. Both auction mechanisms allow the seller to simultaneously auction off two or more units of an item. Bidders must specify in their bid the price per unit and the number of units desired. That is, unlike in the standard formats described above, bidders are not allowed to express a whole demand curve with prices as a function of quantities, but only one price-quantity-pair. Bids are then ranked by price, then by quantity, and finally by the timing of the bid (earlier bids take precedence). There is no proxy bidding. During the auction, bids can be improved according to an improvement rule requiring that the pair value (price times quantity) must increase with any new submitted price-quantity pair. The most

[43] This is similar to market power effects in monopsony. The ranking of the uniform-price auction and the pay-as-bid auction, where bidder too shade their bids, is ambiguous in both efficiency and revenue terms (Engelbrecht-Wiggans and Kahn 1998, Ausubel and Cramton 2002).

important difference between Yankee and Dutch auctions is that in Yankee auctions all winning bidders pay their own bids, while in Dutch auctions all winning bidders pay the same, uniform price, which is the lowest successful bid.[44]

7.3 Complementarities and combinatorial auctions

Auctioning multiple items quickly becomes complicated when there are complementarities between items. Complementarities exist when the value of a bundle of items is larger than the sum of values of each object separately. This is the case in many applications including auctions for the radio spectrum, electricity, airport-landing-slot, supply chains, and transportation services. In such cases, a bidder may end up stuck with items that are worth little because he failed to win complementary items (exposure problem), or he may quit early, fearing that he may fail to win complementary items (holdup problem). As a result, inefficiencies are likely to arise in multi-item auctions where bidders cannot ensure winning complementary items. Theory suggests that in these situations, a combinatorial auction, in which bidders can place bids for one or more *packages* of items, can increase revenue and efficiency. The underlying reason is that these auctions allow bidders to more fully express their preferences. Applications of combinatorial auctions include truckload transportation, bus routes, industrial procurement, airport arrival and departure slots auctions, radio spectrum auctions, and course registration at Chicago Business School (see Cramton et al. forthcoming and Kwasnica et al. 2005 and the references cited therein).

The most famous combinatorial auction is the *Vickrey-Clarke-Groves* (VCR) mechanism, sometimes called generalized Vickrey auction, which works as follows. Bidders bid on all possible packages. The items are then allocated to bidders such that efficiency (the sum of realized values) is maximized according to the stated bids. Each winner pays the smallest (fictitious) bid such that he would still have won his part of the allocation. The resulting price for a bidder equals the *external costs* (social shadow costs) of winning, in the sense that it is the (stated) value of the awarded package for the *other* bidders. Observe that this holds analogously for the single-object second-price auction introduced in section 2 and Vickrey's multi-unit

[44] As long as demand is smaller than supply, the price equals the seller's reservation price.

auction introduced in the last subsection. The VCG mechanism generalizes these formats. In particular, bidding one's values for each package is a dominant strategy.[45]

However, the VCR mechanism suffers from a number of practical problems in the presence of complementarities that seem to seriously limit its usefulness for many applications (Ausubel and Milgrom 2004).[46] One is that the VCG mechanism does not maximize revenues. In fact, revenues can be very low when items are not substitutes though competition is substantial. This alone disqualifies the auction for many purposes.[47] Furthermore, the VCG mechanism makes it easy for losing bidders to collude, and individual bidders can sometimes profit from bidding under pseudonyms – something that appears to be particularly problematic for online auctions, where identities can be more easily manipulated. Another problem is computational complexity. The number of potential bids per bidder is exponentially growing with the number of items auctioned. There are $2^N - 1$ packages of N items. Bidding on all packages can be too demanding for human bidders, even though the VCG mechanism removes all *strategic* complexity by implementing dominant strategies. Computational complexity is also an issue for the auctioneer. Finding the efficiency (or revenue) maximizing allocation of objects in a general combinatorial auction is difficult (more precisely: NP-hard; see de Vries and Vohra 2003), though researchers succeeded in demonstrating that the 'winner determination' problem can often be reasonably addressed (e.g., Rothkopf et al. 1998, Sandholm et al. 2005). Finally, it has been shown that no general equivalent of the VCG mechanism exists in common-value environments, and second-best mechanisms have not yet been identified (Jehiel and Moldovanu 2001).

That said, there has been much progress in *practical* multi-object auction design in recent years. Much of the literature focuses on open, progressive auctions, which can reduce strategic and computational complexity. Some researchers argue, based on applied research in electricity and other infrastructure industry markets, that when complementarities are weak and do not strongly differ across bidders, auction formats like the simultaneous ascending auction may work

[45] Bidders must be risk-neutral and are not allowed to face binding budget constraints. Ausubel (forthcoming) developed an open auction version of the VCG mechanism. Another well-known combinatorial auction is the pay-as-bid package auction by Bernheim and Whinston (1986), which is relatively easy and transparent, but strategically much more complex, as it is typically the case with pay-as-bid auction formats.

[46] Maskin (2005) entertains a more positive picture of the potential practical importance of the VCG mechanism.

[47] General results about revenue maximizing auction mechanisms in the private-value multi-object environment do not exist.

satisfactorily, even though they do not allow combinatorial bids (e.g., Ausubel and Cramton 2004). A well-known progressive auction format that includes combinatorial bidding features is Banks et al.'s (1989) continuous-package bid auction that tries to reduce both value and strategic computation (see Kwasnica et al. 2005 for a recent advancement of this format). Another interesting format to deal with complementarities is the *ascending-proxy auction* (Ausubel, Cramton and Milgrom forthcoming). This is a hybrid auction that begins with a simple and transparent clock phase, not unlike the simultaneous-clock auction, and that ends with a final proxy auction round based on package bids. Similar to eBay's proxy bidding system, bidders in the proxy phase submit values to an artificial proxy agent who then bids on their behalf to maximize profits. It can be shown that including such a proxy phase may handle many of the complications that we discussed above, including the exposure problem.

Some experimental studies, starting with Banks et al. (1989; see also Ledyard et al. 1997, and Plott 1997), investigate bidding when complementarities are present. However, to our knowledge, the only experimental paper that relates its design directly to online auctions is Katok and Roth (2004). They compared the performance of an auction designed to resemble eBay's multi-unit "Dutch" auction to the descending-price auction. The laboratory setting used a set of value environments that include more or less strong complementarities among homogenous objects. Overall, eBay's ascending Dutch auction performed relatively poorly because of the exposure problem. Recall that while the eBay mechanism guarantees a uniform price for all units, it does not guarantee a winning bidder the entire quantity on which he bids. On the other hand, the descending Dutch auction avoids the exposure problem because a bidder who stops the clock obtains the full quantity he desires at the price he stopped the auction. In this sense, the descending Dutch auction can be interpreted as a simple version of a combinatorial auction in case of homogeneous goods. Katok and Roth (2004) conclude that eBay's Dutch auction is susceptible to the exposure problem in environments with synergies, but they also mention that synergies may not be very relevant for B2C and C2C auctions such as eBay. We add that eBay gives winners the right to refuse to purchase 'partial quantities'– a rule that has not been accounted for in the experiment. That is, if a bidder only wins some of the desired object, he does not have to buy any of them. This rule is meant to protect eBay users from the exposure problem (but might create other strategic complications as we will briefly note in the next section).

Section 8: Design of online auctions

Auction design matters. In the previous sections, we have shown that the choice of the auction format, the reservation price, the buy-it-now price and other auction parameters may systematically and significantly affect revenue, efficiency and bidder participation. In this section, we discuss some further auction mechanism choices relevant to online auctions, which have been studied in the literature. [48]

8.1 The advantages of long, open auctions

Unlike offline auctions that typically last only a few minutes, Internet auctions such as those on eBay, Yahoo and Amazon last many days.[49] Since bidders may enter an auction from anywhere, and at anytime, a longer auction time allows more bidders to spot an item and bid on it. Lucking-Reiley et al. (1999) and Hasker et al. (2004) observed that longer auction durations on eBay tend to attract more bidders and lead to higher prices. Lucking-Reiley et al. (1999) reported that 7-day auction prices are approximately 24 percent higher than shorter auctions, and 10-day auctions are 42 percent higher, on average. Hasker et al. (2004) observed that the change in the final sales price achieved by extending the auction from three to ten days is about 10.9 percent.[50]

Long durations also create challenges because bidders cannot be expected to continually monitor the auctions. Many auction houses, including eBay, respond to this by providing bidders with artificial proxy agents. These agents bid on the bidders' behalf, automatically responding as other bids come in, and thus free bidders from the necessity of following the auctions and the price discovery process themselves.

A related important design question is whether the auction should be conducted by sealed-bid. That is, should bidders submit their (proxy) bids over an extended period of time, but without the opportunity to react to the bidding activity of other human or proxy bidders, or should bidding be open, so that bidders can see how bidding activity evolves during the course of

---

[48] For the question how to promote trust and trustworthiness in online auctions by clever design choices, see e.g. Dellarocas (this Handbook), as well as Brosig et al. (2003), Bolton et al. (2004a,b, forthcoming), Güth et al. (2005), among many others.

[49] Google's and Yahoo's auctions of online ads are even always accepting bids.

[50] Hasker et al. (2004) also reported that experienced sellers respond to these incentives in that they sell more valuable objects in longer auctions. Simonsohn (2005) found, on the other hand, that too many sellers set their auctions to end during peak-demand hours such that the probability of sale during such hours is actually lower.

auction? This way, bidders would retain the right to change (proxy) bids in response to the bid history.[51] Many online auction houses, such as eBay, chose an open format. From a theoretical point of view, open ascending-price auctions tend to reduce the force of the 'winner's curse' in environments with a common-value element, because the competitors' bidding activities may convey relevant information that the bidders use in revising their estimates of value. Thus, uncertainty is reduced and so is the winner's curse, and bidders can bid more aggressively. This, in turn, can result in higher revenues in open auctions (see Milgrom and Weber (1982) for the theory and Klemperer (1999) for a more precise intuition behind the effect). Recently, Compte and Jehiel (2004) showed that open auctions are also the preferred choice in private-value environments, if the bidders do not know ones' value a priori. So the rationale for using open formats appear quite robust across auction models.

On the other hand, note that such results, derived in simple auction models, cannot directly be applied to online auctions. One characteristic of eBay is that bidders can enter and leave the auction at any point they wish. So bidding activity, or non-activity, has less information value than in the ascending-price auction described above, in which entry and exit decisions are perfectly observable. Yet, the fact that bidders condition their behavior on others' activities (see, e.g., sections 3 and 4) suggests that open online auctions reveal some helpful information.

Another argument for open auction formats comes from laboratory experiments. It has been shown that the feedback delivered in open second-price auctions such as eBay substantially accelerates the speed of learning compared to second-price sealed-bid auctions (Ariely et al. forthcoming). This improves the price discovery process and increases competition among bidders so that efficiency and revenues can be enhanced, even in purely private-value environments. In line with this finding, Ivanova-Stenzel and Salmon (2004) report that, when having the choice between sealed-bid and open, ascending-bid auctions, laboratory subjects in a private-value environment have a strong preference for the open format. Finally, Cramton (1998) notes that in practical applications, the dynamic price-discovery process of an open auction most often does a better job than sealed bidding. This is, of course, in particular true in multi-object auctions, where the dynamics facilitate arbitrage and packaging.

[51] We assume here that the open auction is ascending.

However, there are also disadvantages that come with open bidding. Open auctions are more susceptible to various forms of collusion and fraud. Bapna (2003) argues that open auctions facilitate collusive bidding against a repeat seller (and has other more technological disadvantages). He therefore recommends that eBay run sealed-bid auctions. The literature on spectrum auctions, however, demonstrated that certain auction design features can address and mitigate many of these problems. For instance, a history of bids that conceals bidder identities can, to some extent, suppress bidder collusion against sellers and rival bidders.[52] Furthermore, the anonymity and the number of potential bidders, as well as free entry in online auctions, seem to make coordination, signaling and communication among bidders more difficult than in many offline auction environments.

Other concerns are probably more justified: Open auctions can lead to lower revenues when bidders are risk-averse (as we mentioned in section 2), and when *ex-ante* asymmetries among bidders are strong or competition is weak (e.g., Cramton 1998). This might be part of the reason why eBay recently introduced a sealed-bid format as an option for sellers; in the '*best offer*' format, bidders can make sealed-bids, and sellers can accept any bid at any time they wish.

8.2. Controlling the pace of bidding

As we have seen in section 4, bidders in eBay auctions tend to bid late. This may distort the virtues of long, open auctions described above. One way to avoid late bidding and to control the pace of auctions is to create pressure on bidders to bid actively from the start. Milgrom and Wilson designed an activity rule that was applied to the U.S. spectrum auctions (McAfee and McMillan 1996). The activity rule requires a bidder to be 'active' (that is to be the current high bidder or to submit new bids) on a predetermined number of spectrum licenses. If a bidder falls short of the required activity level, the number of licenses it is eligible to buy shrinks. Thus, bidders are prevented from holding back. However, activity rules of this sort are incompatible with the flexibility needed on global auction platforms.

Roth and Ockenfels (2002) observed that the rule by which online auctions end may have a substantial effect on the timing of bids and price discovery. On eBay, auctions end at a predetermined time: a 'hard close'. In contrast, Amazon emulates the 'Going, Going, Gone'

[52] See Klemperer (2004) for a review of design recommendations for the European spectrum auctions to avoid collusive behavior.

feature of traditional auction houses. That is, Amazon automatically extends an auction if a bid comes in late, so that all bidders always have the opportunity to respond to the opponents' bids.[53]

Ockenfels and Roth (forthcoming) show that, although the risks of last-minute bidding remain, the strategic advantages of last-minute bidding are eliminated or severely attenuated in Amazon-style auctions. That is, a bidder who waits to bid until the last seconds of the auction still runs the risk that his bid will not successfully be transmitted in time. However, if his bid is successfully transmitted, the auction will be extended for ten minutes, so that, no matter how late the bid was placed, other bidders will have time to respond. Thus on Amazon, an attentive incremental bidder, for example, can respond whenever a bid is placed.[54] The differences in the strategic environment are reflected in the data of Roth and Ockenfels (2002): There is significantly more late bidding on eBay than on Amazon. For instance, 40 percent of eBay-computer auctions and 59 percent of eBay-antiques auctions in the sample have last bids in the closing 5 minutes, compared to about 3 percent of both Amazon computer and Amazon antiques auctions that have last bids in the final five minutes before the initially scheduled deadline or later. Further analysis reveals that while the impact of the bidders' feedback numbers on late bidding is significantly positive in eBay, it is negative in Amazon, suggesting that more experienced bidders on eBay bid later than less experienced bidders, but experience in Amazon has the opposite effect.

Experiments by Ariely et al. (forthcoming) replicate these findings in a controlled laboratory private-value setting in which the only difference between auctions is the ending rule. The experiment thus controls for differences other than the closing rule that might affect behavior on Amazon and eBay, such as the number of auctions being conducted at a time and the number of potential bidders. The experiment also demonstrates that, ceteris paribus, 'early' prices on Amazon are an increasingly good predictor for final prices, whereas price discovery on eBay

[53] In Amazon's own words: "We know that bidding can get hot and heavy near the end of many auctions. Our Going, Going, Gone feature ensures that you always have an opportunity to challenge last-second bids. Here's how it works: Whenever a bid is cast in the last 10 minutes of an auction, the auction is automatically extended for an additional 10 minutes from the time of the latest bid. This ensures that an auction can't close until 10 minutes have passed with no further bids." On Yahoo's auction platform, the seller decides whether he wishes a hard or a soft close. Otherwise, all three platforms employ similar auction rules.

[54] However, there are other, non-strategic reasons for late bidding, including procrastination, use of search engines that make it easy to find auctions about to end, endowment effects, or management of bidding in multiple auctions in which similar objects may be offered. These motives for late bidding should be relatively unaffected by the difference in closing rules between eBay and Amazon.

became increasingly frenzied. Simulation experiments by Duffy and Ünver (2005) with artificial adaptive agents who can update their strategies via a genetic algorithm, replicate these findings and thus provide another robustness check.

Controlled field experiments, on the other hand, seem to have more difficulties finding evidence for the impact of the ending rule. Brown and Morgan (2005) and Houser and Wooders (2005) took advantage of the fact that Yahoo sellers are allowed to choose whether to end the auction with a hard or a soft close. In both studies, identical items were sold using both ending rules. None of these studies found a significant effect of the ending rule on the amount of late bidding.[55] However, Houser and Wooders (2005) observed – as Ariely et al. (forthcoming) and Duffy and Ütku (2005) – that, ceteris paribus, hard-close auctions tend to raise less revenue than soft-close auctions.[56]

Online market design also includes the design of artificial agents, such as eBay's 'proxy bidder'. Because late bidding involves a good deal of planning and effort, artificial agents can also help executing late-bidding strategies. In fact, there is a market for artificial sniping agents that will allow a bidder not only to submit a proxy bid, but also to do so at the last moment. Sniping agents take two forms: downloadable programs that run on the bidder's own computer, and web-based services like esnipe.com to which a bidder can subscribe. Both offer bidders the ability to choose their maximum bid early in the auction, record when the auction is scheduled to end, and decide how many minutes or seconds before the end of the auction the sniping agent should submit the bid.

Recall that whether the timing of bids matters depends on the rules of the game. Artificial last-minute bidding agents (like esnipe.com) might support human bidders in eBay auctions, but they would hardly help on Amazon, where the closing rule removes or greatly attenuates the incentives to snipe. By the same token, human bidders on Amazon have more reason to make use of the proxy bidding agent provided by the auction houses than bidders on eBay, where the fixed deadline may create incentives to submit the bid late, depending on other (human or

[55] In a recent laboratory experiment, in which three sudden termination variants of hard-close auction (a.k.a. *candle auction*) were examined, Füllbrunn and Sadrieh (2006) find that the extent of late-bidding crucially depends on the first stage in which the probability of sudden termination is greater than zero.

[56] In a theoretic model of sequential auctions, Stryszowska (2005b) identified a situation in which soft-close auctions should be expected to yield smaller revenues.

artificial) bidders' behavior. Thus, how well different kinds of artificial agents perform depends on how the auction rules are designed.

Note also, that as sniping by human and artificial agents become more widespread on eBay, eBay will be gradually transformed into a sealed-bid second-price auction. If a large part of the late bidding activity takes place on third party sites like esnipe.com, eBay faces a number of design and rule choices; one is to ban sniping services. In fact, eBay.de (Germany) banned third party sniping services in its general terms and conditions (which is, of course, difficult to enforce), because, according to them, bidders who use sniping services have an "unfair advantage" over people who bid manually. A second choice would be just the opposite: recapturing the sniping market by offering a sniping option on eBay itself. Under this option, last minute bids submitted in advance directly to eBay could all be counted at the same time, immediately after the auction close. This would give bidders certainty, both that their bids would be successfully transmitted, and that there would be no time for other bidders to react. Of course, if all bidders used this option, the auction becomes a sealed-bid auction (Ockenfels and Roth 2002). As we have argued above, eBay might prefer not to encourage this development towards sealed-bids, given the advantages of open auctions, yet even now, eBay is considering a sniping service that would enable last-minute bidding via phone (http://www.unwiredbuyer.com/). While bidding by phone will still involve risks that the bid fails to successfully register, it will likely further increase the number of snipes. Finally, eBay could consider changing the ending rule of the auction to a soft close. This, however, may also cause adverse effects such as lowering the entertainment value of eBay.[57]

There is at least one other design choice influencing auction design: the speed of the price clock in decreasing-price auctions. As we mentioned in section 2, Lucking-Reiley (1999) found

---

[57] Ockenfels (2003) noted that online negotiation sites that promise dispute resolution (such as e-commerce disputes and traditional litigation) via electronic and standardized communication have to deal with related design problems. One of the more prominent online negotiation sites, clicknsettle.com, experimented in 1999 with round-by-round demands and offers. But this format did not prove to be effective, because a deadline effect similar to what has been observed on eBay and to what has been observed in experimental bargaining games (Roth et al. 1988) hindered efficient negotiations: "*After reviewing the early results with our clients, we discovered that in most negotiations, the first two rounds were being 'wasted' and the disputing parties really only had one opportunity to settle the case, the final round.*" (http://www.clicknsettle.com/onlinenegmodel.cfm 2003). EBay also works together with a dispute resolution provider. A recent study by Brett et al. (2005) investigated the time it takes to resolve a dispute in an on-line setting. By analysing 582 eBay-generated disputes they find that the opening moves can be critical in accelerating or delaying resolution to disputants.

that in a controlled field study of online auctions, the single-item descending-price format yields higher revenues than corresponding sealed-bid auctions – just the opposite of what has been found in some laboratory experiments (Cox et al. 1982, 1983). Lucking-Reiley observed that his descending-price auctions took much longer than the experiments and speculated that the higher revenues are because bidders may be impatient to complete their purchase, or having more time allows bidders to more accurately determine their value of the item. These ideas are supported in a laboratory experiment by Katok and Kwasnica (2004). In descending-price auctions with a fast clock, revenue turned out to be significantly lower, and with a slow clock significantly higher, than in the corresponding sealed-bid version. As the authors show, this bidding pattern is in line with a simple model of impatient bidders. Carare and Rothkopf (2005) come to similar conclusions in both a decision-theoretic and game-theoretic model that incorporates a 'cost of returning to the auction site'. Bidders prefer to bid sooner, yielding higher prices, when the cost is higher. These results suggest that, without seller competition and without impatience on the side of the seller, sellers would prefer to implement a slow clock rather than a fast clock, or a sealed-bid mechanism. We are not aware of any online auction that allows a choice like this, but eBay's buy-it-now option may be a substitute choice, because it gives impatient bidders the opportunity to immediately end the auction at a higher price.

8.3 Design aspects in multi-unit auctions

There are potentially many ways to sell multiple items in online auctions. One way is to sell them simultaneously or sequentially in a series of single-object auctions. Another way is to sell them through a single auction, tailored to selling multiple units. The latter approach has several advantages. For one, selling multiple units through one auction can reduce transaction costs for both buyers and sellers. Second, simultaneous and sequential auctions impose strategic complexities and coordination problems on bidders, because bidders must guess prices of the other objects in order to realize arbitrage and to efficiently package objects. Wrong guessing may hamper revenue and efficiency (Milgrom 2004, Cramton 2004).

Laboratory and field studies of sequential auctions strongly support this view. For instance, sequential auctions typically fail to generate similar prices for similar items. Rather, prices display a downward drift. This phenomenon is called '*declining price anomaly*' (Ashenfelter 1989). One possible explanation is related to the winner's curse: those bidders who win the early units are those who overestimated the prices realized in later auctions. Related problems arise in simultaneous eBay auctions. Since eBay bidders have only weak incentives to contribute to the price discovery process early in the auction – especially in the presence of simultaneous, multiple listings of identical items – the decision where and what to bid is complex and may lead to random or erroneous entry decisions close to the ending time. The result is coordination failure. Stryszowska (2005a,b,c) investigated coordination (failure) in simultaneous, overlapping and sequential Internet auctions.

In principle, multi-unit auction formats, such as eBay's "Dutch" auction, diminish coordination problems by reducing the number of auctions. They also enforce uniform prices for identical objects, reducing the risk associated with price dispersion. However, details in the design matter. Ockenfels (2005) noted that even in the simplest case of unit-demand, the price rule of eBay's "Dutch" auction makes bidding much more complicated than in the single-object auction. Assuming that no bidder demands more than one unit, the 'natural' extension of eBay's second-price single-object auction mechanism is the Vickrey auction described earlier, in which the final price is equal to the *highest losing bid* (plus a small increment). Facing Vickrey's price rule, bidders should 'sooner or later' just bid their values, independent of the other bidders' behavior. In eBay's "Dutch" auction, however, the final price is equal to the *lowest winning bid*, so that one of the winners will eventually determine the price. This creates incentives for bid shading. Specifically, winners can minimize the price paid by not bidding more than a small increment above the highest losing bid.[58] But, because the highest losing bid is usually not known before the auction is over, the outcome of the auction again depends on the accurateness of the bidders' estimations. In fact, Ockenfels (2005) found more bid shading in eBay's "Dutch" auction than in eBay's single-object auction in a controlled field experiment on eBay.

These kinds of arguments convinced eBay Germany to change the multi-unit format in the summer of 2005. The multi-unit auction format is now called "multi-auction" and to a large extent is analogous to the single-object auction. Most importantly, there is now proxy bidding (proxy bids are concealed to other bidders), and the price equals the highest losing bid – analogous to the single-object format. However, there are other issues with eBay's multi-unit auction having to do with the fact that neither the old, nor the new format avoid demand reduction and exposure problems for multi-unit demand. We believe that designing a robust

[58] In the single-object auction, finding this price is the job of the proxy bidder.

multi-unit auction that takes complex preferences and incentives of the bidders into account is still an important challenge for online auctions. The slight changes in Germany are moving in the right direction.

References

Anderson, S. T., D. Friedman, G. H. Milam, and N. Singh (2004). Buy it Now: A Hybrid Internet Market Institution. UC Santa Cruz Economics Working Paper No. 565. http://ssrn.com/abstract=533182.

Anwar, S., R. McMillan, M. Zheng (2006). Bidding Behavior in Competing Auctions: Evidence from eBay. *European Economic Review* 50(2). 307-322.

Ariely, D., A. Ockenfels, and A.E. Roth (forthcoming). An Experimental Analysis of Ending Rules in Internet Auctions. *The RAND Journal of Economics*.

Ariely, D. and I. Simonson (2003). Buying, Bidding, Playing or Competing?: Value Assessment and Decision Dynamics in Online Auctions. *Journal of Consumer Psychology* 13(1–2). 113–123.

Arora, A., H. Xu, R. Padman, and William Vogt (2003). Optimal Bidding in Sequential Online Auctions. Working Paper #2003-4. Carnegie Mellon University. http://www.heinz.cmu.edu/wpapers/author.jsp?id=ashish.

Ashenfelter, O. (1989). How Auctions Work for Wine and Art. *Journal of Economic Perspectives* 3. 23-36.

Asker, J., B. Grosskopf, C. N. McKinney, M. Niederle, A. E. Roth, and G.Weizsäcker (2004). Teaching Auction Strategy using Experiments Administered via the Internet. *Journal of Economic Education* 35(4). 330 – 342.

Ausubel, L.M. (forthcoming). An Efficient Ascending-Bid Auction for Multiple Objects. *American Economic Review*.

Ausubel, L.M. and P. Cramton (2002). Demand Reduction and Inefficiency in Multi-unit Auctions. Working paper. University of Maryland.

Ausubel, L. M. and P. Cramton (2004). Auctioning Many Divisible Goods. *Journal of the European Economic Association* 2. 480-493.

Ausubel, L. M., P. Cramton and P. Milgrom (forthcoming). The Clock-Proxy Auction: A Practical Combinatorial Auction Design. In: P. Cramton, R. Steinberg and Y. Shoham (eds.). *Combinatorial Auctions*.

Ausubel, L. M., and P. Milgrom (forthcoming). The Lovely but Lonely Vickrey Auction. In: P. Cramton, R. Steinberg and Y. Shoham (eds.). *Combinatorial Auctions*.

Bajari, P. and A. Hortaçsu (2003). Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions. *Rand Journal of Economics* 2. 329-355.

Bajari, P. and A. Hortaçsu (2004). Economic Insights from Internet Auctions. *Journal of Economic Literature*, XLII, 457-486.

Bajari, P. and A. Hortaçsu (2003). Economic Insights from Internet Auctions: A Survey. NBER Working Paper No. W10076. http://ssrn.com/abstract=467544.

Ball, S., M.H. Bazerman, and J.S. Carroll (1991). An Evaluation of Learning in the Bilateral Winner's Curse. *Organizational Behavior and Human Decision Processes* 48. 1-22.

Banks, J., J. Ledyard, and D. Porter (1989). Allocating uncertain and unresponsive resources: An experimental approach. *RAND Journal of Economics* 20(1). 125.

Barbaro, S. and B. Bracht (2004). Shilling, Squeezing, Sniping: Explaining Late Bidding in Online Second-price Auctions. Working paper. University of Mainz, Germany.

Bapna, R. (2003). When Snipers become Predators: Can Mechanism Design Save Online Auctions? *Communication of the ACM*, 46(12), 152-158.

Bapna, R., P. Goes, and A. Gupta (2000). A theoretical and empirical investigation of multi-item on-line auctions. *Information Technology and Management* 1(1). 1–23.

Bapna, R., P. Goes, and A. Gupta (2001). Insights and Analyses of Online Auctions. *Communications of the ACM* 44(11). 42-50.

Bernheim, D. and M. Whinston (1986). Menu Auctions, Resource Allocation and Economic Influence. *Quarterly Journal of Economics* 101. 1-31.

Bolton, G. and A. Ockenfels (2000). ERC – A theory of equity, reciprocity and competition. *American Economic Review* 90(1). 166-193.

Bolton, G., and A. Ockenfels (2006). Bridging Laboratory and Field Behavior: Bargaining and Auctions on EBay. Working paper. University of Cologne.

Bolton, G. E., E. Katok, and A. Ockenfels (2004). How Effective are Online Reputation Mechanisms? An Experimental Investigation. *Management Science* 50(11). 1587-1602.

Bolton, G. E., E.Katok, and A. Ockenfels (2004). Trust among Internet Traders. A Behavioral Economics Approach. *Analyse und Kritik* 26. 185-202.

Bolton, G. E, E. Katok, and Axel Ockenfels (forthcoming). Bridging the Trust Gap in Electronic Markets. A Strategic Framework for Empirical Study. In: E Akcali, J. Geunes, P.M. Pardalos, H.E. Romeijn and Z.J. Shen (eds.). *Applications of Supply Chain Management and E-Commerce Research in Industry*. Kluwer Academic Publishers.

Bose S. and A. Daripa (2006). Buy-Now Option in Auctions. Working Paper. University of Illinois, Urbana-Champaign.

Borenstein, S., J. Bushnell, and Frank Wolak (2002). Measuring Market Inefficiencies in California's Wholesale Electricity Industry. *American Economic Review* 92 (5). 1376-1405.

Börgers, T. and E. van Damme (2004). Auction Theory for Auction Design. In: M. Janssen (ed.). Auctioning Public Assets: Analysis and Alternatives. Cambridge University Press. 19-63.

Brett, J. M., M. Olekalins, N. Goates, R. A. Friedman, and C. C. Lisco (2005). Tortoise or Hare? A Study of the Speed to Resolution in On-Line Disputes. Working paper. Kellogg School of Management, Northwestern University.

Brosig, J., A. Ockenfels, and J. Weimann (2003). The Effect of Communication Media on Cooperation. *German Economic Review* 4(2). 217-241.

Brosig, J., and J.P. Reiss (forthcoming). Entry decisions and bidding behavior in sequential first-price procurement auctions: An experimental study. *Games and Economic Behavior*.

Brown, J. and J. Morgan (2005). How much is a Dollar Worth? Tipping versus Equilibrium Coexistence on Competing Online Auction Sites. Working paper. University of California, Berkeley.

Budish, E. and L. Takeyama (2001). Buy prices in online auctions: irrationality on the Internet? *Economics Letters* 73. 325-333.

Burguet, R. and J. Sákovics (1999). Imperfect Competition in Auction Design. *International Economic Review* 40(1). 231-247.

Bulow, J.I. and P.D. Klemperer (1994). Rational Frenzies and Crashes. *Journal of Political Economy* 102. 1-23.

Cabral, L. and A. Hortaçsu (2005). The Dynamics of Seller Reputation: Theory and Evidence from eBay. Working Paper. New York University. http://pages.stern.nyu.edu/~lcabral/workingpapers/CabralHortacsu_June05.pdf.

Camerer, C. F. (2003). Behavioral Game Theory: Experiments in Strategic Interaction. Princeton University Press.

Carare, O. and M. Rothkopf (2005). Slow Dutch auctions. *Management Science* 51(3). 365-373.

Cassady, R. 1967. Auctions and Auctioneering. California University Press, California.

Chakraborty, I. and G. Kosmopoulou (2004). Auctions with shill bidding. *Economic Theory*. Springer. vol. 24(2). 271-287.

Chen, R., G. Janakiraman, R. Roundy, and R. Q. Zhang (2005). Efficient auction mechanisms for supply chain procurement. *Management Science* 51(3). 467-482.

Cox, J. C. and R. Mark Isaac (1984). In Search of the Winner's Curse. *Economic Inquiry*. vol. 22. 579-592.

Cox, J., B. Roberson, and Vernon L. Smith (1982). Theory and Behavior of Single Object Auctions. *Research in Experimental Economics* 1. 61-99.

Cox, James, Vernon L. Smith and James M. Walker (1983). A Test that Discriminates Between Two Models of the Dutch-First Auction Non-Isomorphism. *Journal of Economic Behavior and Organization* 4. 205-219.

Cox, J. C., V.L. Smith, and J.M. Walker (1988). Theory and Individual Behavior of First Price Auctions. *Journal of Risk and Uncertainty*. vol. 1, pp. 61-99.

Cramton, P. (1995). Money Out of Thin Air: The Nationwide Narrowband PCS Auction. *Journal of Economics and Management Strategy* 4. 267-343.

Cramton, P. (1998). Ascending Auctions. *European Economic Review* 42(3-5). 745-756.

Cramton, P. (2002). Spectrum Auctions. In: M. Cave, S. Majumdar, and I. Vogelsang (eds.). *Handbook of Telecommunications Economics*. Amsterdam: Elsevier Science B.V. 605-639.

Cramton, P. (2004). Simultaneous Ascending Auctions. Working paper. University of Maryland.

Cramton, P., Y. Shoham, and R. Steinberg (forthcoming). Introduction to Combinatorial Auctions. In: P. Cramton, Y. Shoham, and R. Steinberg (eds.). *Combinatorial Auctions*. MIT Press.

Compte, O. and P. Jehiel (2004). Auctions and Information Acquisition: Sealed-bid or Dynamic Formats? Working paper. C.N.R.S. Paris.

Dellarocas, C. (forthcoming). The Digitization of Word-of-Mouth: Promise and Challenges of Online Reputation Mechanisms. *Management Science*.

Deltas, G., R. Engelbrecht-Wiggans (2005). Naive bidding. *Management Science* 51(3). 328-338.

de Vries, S. and R. Vohra (2003). Combinatorial Auctions: A Survey. *INFORMS J. Comput* 5(3). 284-309.

Dennehy, M. (2000). eBay Adds 'Buy-It-Now' Feature. *AuctionWatch.com* October 26:2000.

Dewally, M. and L. H. Ederington (2004). What Attracts Bidders to Online Auctions and What is Their Incremental Price Impact? http://ssrn.com/abstract=589861.

Dholakia, U. M. and Kerry Soltysinski (2001). Coveted or Overlooked? The Psychology of Bidding for Comparable Listings in Digital Auctions. *Marketing Letters* 12(3). S. 225-237.

Dholakia, U. M., S. Basuroy, and K. Soltysinski (2002). Auction or agent (or both)? A study of moderators of the herding bias in digital auctions. *International Journal of Research in Marketing* 19. S. 115-130.

Ding, M., J. Eliashberg, J. Huber, and R. Saini (2005). Emotional bidders: An analytical and experimental examination of consumers behavior in priceline-like reverse auctions. *Management Science* 51(3). 352-364.

Duffy, J. and U. Ünver (2005). Internet Auctions with Artificial Adaptive Agents: A Study on Market Design. Working paper. University of Pittsburgh.

Durham, R. and Standifird (2004). eBay's Buy-It-Now Function: Who, When, and How. *Topics in Economic Analysis & Policy*. vol. 4. issue 1. 1316-1316.

Easley, R. F. and R. Tenorio (1999). Bidding Strategies in Internet Yankee Auctions.

Edelman, B., M. Ostrovsky, and Michael Schwarz (2005). Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. Working paper. Harvard University.

Elyakime, B., J.J. Laffont, P.Loisel, and Q.Vuong (1994). First-Price Sealed-Bid Auctions with Secret Reservation Prices. *Annales D'Economie et de Statistique*. Vol. 34. 115-141.

Engelberg, J. and J. Williams (2005). A License to Shill: Incremental Bidding and eBay's Proxy System. Working paper. Kellogg School of Management.

Engelberg, J. and J. Williams (2005). EBay's Proxy System: A License to Shill. Available at SSRN: http://ssrn.com/abstract=616782

Engelbrecht-Wiggans, R. (1987). On Optimal Reservation Prices in Auctions. *Management Science* 33 (6), 763-770.

Engelbrecht-Wiggans, R. (1989). The Effect of Regret on Optimal Bidding in Auctions. *Management Science* 35(6). 685-692.

Engelbrecht-Wiggans, R. and T. Nonnenmacher (1999). A Theoretical Basis for 19th-Century Changes to the Port of New York Imported Goods Auction. *Explorations in Economic History*. Elsevier. vol. 36(3). 232-245.

Engelbrecht-Wiggans, R., E. Haruvy, and E. Katok (2005). Market Design for Procurement: Empirical and theoretical investigation of buyer-determined multi-attribute mechanisms. Working Paper. http://lema.smeal.psu.edu/katok/ma.pdf.

Engelbrecht-Wiggans, R. and Charles M. Kahn (1998). Multi-unit Auctions with Uniform Prices. *Economic Theory* 12. 227-258.

Engelbrecht-Wiggans, R., C. Kahn (2005). Low-revenue equilibria in simultaneous ascending-bid auctions. *Management Science* 51(3). 508-518.

Engelbrecht-Wiggans, R. and E. Katok (2005a). Regret in Auctions: Theory and Evidence. Working Paper. http://lema.smeal.psu.edu/katok/regretTheory_post.pdf.

Engelbrecht-Wiggans, R. and E. Katok (2005b). A Tough Test of Risk Aversion in First Price Sealed Bid Auctions. Working Paper. http://lema.smeal.psu.edu/katok/riskAversion_post.pdf.

Engelbrecht-Wiggans, R., J. A. List, and David H. Reiley (forthcoming). Demand Reduction in Multi-unit Auctions with Varying Numbers of Bidders: Theory and Field Experiments. *International Economic Review*.

Engelmann, D. and V. Grimm (2004). Bidding Behavior in Multi-Unit Auctions - An Experimental Investigation and some Theoretical Insights. http://home.cerge-ei.cz/engelmann/WorkingPapers/DerExGEB.pdf.

Etzion, H., E. Pinker, and A. Seidmann (2003). Analyzing the simultaneous use of auctions and posted prices for online selling. Working paper CIS-03-01. University of Rochester.

Fehr, E.,and K. Schmidt (1999). A Theory of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics* 114. 817-868.

Füllbrunn, S. and A. Sadrieh (2006). Sudden Termination Auctions. Working Paper. University of Magdeburg.

Graham, D. A., R. C. Marshall, and Jean-Francois Richard (1990). Phantom Bidding Against Heterogeneous Bidders. *Economics Letters* 32. 13-17.

Gray, S. and D. Reiley (2004). Measuring the Benefits to Sniping on eBay: Evidence from a Field Experiment. Working paper. University of Arizona.

Grimm, V., F. Riedel, and E. Wolfstetter (2003). Low Price Equilibrium in Multi–Unit Auctions: The GSM Spectrum Auction in Germany. *International Journal of Industrial Organization*.21. 1557-1569.

Grosskopf, B., Y. Bereby-Meyer, and M. H. Bazerman (2003). On the Robustness of the Winner's Curse Phenomenon. Harvard Business School. Working Paper No. 03119.

Güth, W., F. Mengel, and A. Ockenfels (2005). The Dynamics of Trust and Trustworthiness on EBay. An Evolutionary Analysis of Buyer Insurance and Seller Reputation. Working paper. University of Cologne.

Gupt, S. and J. Gallien (2004). Temporary and Permanent Buyout Prices in Online Auctions. http://hdl.handle.net/1721.1/7468.

Haile, P.A. (2000). Partial Pooling at the Reserve Price in Auctions with Resale Opportunities. *Games and Economic Behavior*, 22, 231-48.

Hasker, K., R. Gonzales, and R. C. Sickles (2004). An Analysis of Strategic Behavior and Consumer Surplus in eBay Auctions. Working Paper. Rice University. http://www.ruf.rice.edu/~rsickles/paper/auction.pdf.

Hayne, S. C., C.A.P. Smith, and L. R. Vijayasarathy (2003a). Sniping in eBay: A Cultural Analysis. Working Paper. University of Colorado.

Hayne, S. C., C.A.P. Smith and L. R. Vijayasarathy (2003b). Who Wins on eBay: An Analysis of Bidders and Their Bid Behaviours. *Electronic Markets* 13 (2003), 4. S.282-293.

Häubl, G., P. T.L. Popkowski Leszczyc (2003). Minimum Prices and Product Valuations in Auctions. *Marketing Science Institute Reports*. Issue 3, No. 03-117. 115-141.

Häubl, G., P. T.L. Popkowski Leszczyc (2004). Bidding Frenzy: Intensity of Competitive Interaction Among Bidders and Product Valuation in Auctions. *Advances in Consumer Research*. Volume 31. 90-93.

Hendricks, K., I. Onur, and T. Wiseman (2005). Preemption and Delay in eBay Auctions. Working Paper. University of Texas at Austin.

Herschlag, M. and R. Zwick. 2000. Internet Auctions: A Popular and Professional Literature Review. *Quarterly Journal of Electronic Commerce* 1(2). 161‐186.

Heyman, J. E., Y. Orhun, and D. Ariely (2004). Auction Fever: The Effect of Opponents and Quasi-Endowment on Product Valuations. *Journal of Interactive Marketing* 18(4), ??.

Hian, C. C. (2003/2004). An Examination of Permanent Buy-now Auction on the Internet.

Hidvegi, Z., W. Wang, and A. B. Whinston (2004). Buy-price English Auction. *Journal of Economic Theory.*

Hof, R. (2001). The People's Company. *Businessweek Online*.

Hoppe, T. and A. Sadrieh (2005). An Experimental Assessment of Confederate Reserve Price Bids in Online Auctions. Working paper. University of Magdeburg.

Hossain, T. (2006). Learning by Bidding. Working paper. Hong Kong University of Science and Technology.

Hossain, T. and J. Morgan (2004). A Test of Revenue Equivalence Theorem using Field Experiments on eBay. Working Paper. University of California at Berkeley. http://faculty.haas.berkeley.edu/rjmorgan/RET%20v7.pdf.

Hossain, T. and J. Morgan (2006). ...Plus Shipping and Handling: Revenue (Non)Equivalence in Field Experiments on eBay. *Advances in Economic Analysis & Policy*.

Houser, D. and J. Wooders (2005). Hard and Soft Closes: A Field Experiment on Auction Closing Rules. In: R. Zwick and A. Rapoport (eds.). *Experimental Business Research*. Vol. II. Kluwer Academic Publishers. 279-287.

Ivanova-Stenzel, R. and T. Salmon (2004). Bidder Preferences among Auction Institutions. *Economic Inquiry* 42. 223-236.

Izmalkov, S. (2004). Shill bidding and optimal auctions. Working Paper. MIT.

Jehiel, P. and B. Moldovanu (2001). Efficient Design with Interdependent Values. *Econometrica* 69. 1237-1259.

Jin, G. Z. and A. Kato (2004). Consumer Frauds and the Uninformed: Evidence from An Online Field Experiment. Working Paper. University of Maryland. http://www.glue.umd.edu/~ginger/research/eBay-Jin-Kato-0305.pdf.

Jin, G. Z. and A. Kato (2005). Quality and Reputation: Evidence from An Online Field Experiment. Working Paper. University of Maryland.

Kagel, J. H. (1995) Auctions: A Survey of Experimental Research. In: J. H. Kagel & A. E. Roth (Eds.). *The Handbook of Experimental Economics*. 501-585. Princeton University Press. Princeton, NJ.

Kagel, J. H. and D. Levin (2001). The Winner's Curse and Public Information in Common Value Auctions. *American Economic Review* 76. 849-920.

Kagel, J. and Dan Levin (1993). Independent Private Value Auctions: Bidder Behavior, in First-, Second- and Third Price Auctions with Varying Numbers of Bidders. *Economic Journal*. July 1993. 103. 868-879.

Kagel, J. H. and D. Levin (1995). Comparative Static Effects of Number of Bidders and Public Information on Behavior in Second-price Common Value Auctions. *International Journal of Game Theory* 24. 293-319.

Kagel, J. H. and D. Levin (2001). Behavior in Multi-unit Demand Auctions: Experiments with Uniform Price and Dynamic Auctions. *Econometrica* 69 (2). 413-454.

Kagel, J. and D. Levin (2002). Common Value Auctions and the Winner's Curse. Princeton University Press.

Katar, R. and D. H. Reiley (2005). Public Versus Secret Reserve Prices in eBay Auctions: Results from a Pokémon Field Experiment. Working paper. University of Arizona.

Katok, E. and A. M. Kwasnica (2004). Time is Money: The Effect of Clock Speed on Seller's Revenue in Dutch Auctions. Working paper. Penn State University.

Katok, E. and A. E. Roth (2004). Auctions of Homogenous Goods with Increasing Returns: Experimental Comparison of Alternative "Dutch" Auctions. *Management Science* 50(8).1044-1063.

Katzman, B. and T. Matthews (2006). The Role of Varying Risk Attitudes in an Auction with a Buyout Option. *Economic Theory* 27. 597-613

Kauffman, R. J. and Charles A. Wood (2003). Running up the Bid: Detecting, Predicting, and Preventing Reserve Price Shilling in Online Auctions. Working paper. MIS Research Center. Carlson School of Management. University of Minnesota. http://misrc.umn.edu/workingpapers/fullPapers/2003/0304_022003.pdf

Kauffman, R. J. and C. A. Wood (2004). Doing Their Bidding: An Empirical Examinations of Factors that Affect a Buyer's Utility in Internet Auctions. http://www.misrc.umn.edu/workingpapers/fullPapers/2002/0204_041903.pdf

Kirchkamp, O. and P. Reiss (2004). The Overbidding-Myth and the Underbidding-Bias. SFB 504 Working Paper No. 04-32. University of Mannheim. http://www.kirchkamp.de/pdf/04-32.pdf

Kirkegaard, R. (2005). "The Phantom Made Me Jump! Preemptive Jump Bidding in (Takeover) Auctions." Working paper, Brock University, October.

Kirkegaard, R. and P.B. Overgaard (2003). Buy-Out Prices in Online Auctions: Multi-Unit Demand. mimeo, University of Aarhus.

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

Kosmopoulou, G. and D. G. De Silva (2005). The Effect of Shill Bidding upon Prices: Experimental Evidence. Available at SSRN: http://ssrn.com/abstract=643143.

Kirchkamp, O., E. Poen, and J.P. Reiss (2004): Outside options: Another reason to choose the first-price auction. Working paper. University of Magdeburg

Krishna, V. (2002). Auction Theory. Academic Press.

Ku, G., D. Malhotra, and J. Keith Murnighan (2005). Towards a Competitive Arousal Model of Decision-Making: A Study of Auction Fever in Live and Internet Auctions. *Organizational Behavior and Human Decision Processes* 96. 89-103.

Kwasnica, A., J. Ledyard, D. Porter, and C. DeMartini (2005). A new and improved design for multi-object iterative auctions. *Management Science* 51(3). 419-434.

Kwon, R., G. Anandalingam, and L. H. Ungar (2005). Iterative combinatorial auctions with bidder-determined combinations. *Management Science* 51(3). 407-418.

Levin, D., and J. L. Smith (1996). Optimal Reservation Prices in Auctions. *Economic Journal* 106, 1271-1282.

Li, H. and G. Tan (2000). Hidden Reserve Prices with Risk Averse Bidders. University of British Columbia. Working Paper.

List, J. (2003). Does Market Experience Eliminate Market Anomalies? *Quarterly Journal of Economics* 118(1). 41-71.

List, J. A. and D. Lucking-Reiley (2000). Demand Reduction in Multi-Unit Auctions: Evidence from a Sportscard Field Experiment. *American Economic Review* 90(4). 961-972.

List, J. A. and D. Lucking-Reiley (2002). Bidding Behavior and Decision Costs in Field Experiments. Economic Inquiry.October 2002. vol. 40, no. 44. 611-619.

Liu, D., J. Chen, and A. B. Whinston (2005). Designing Online Auctions with Past Performance Information. Working paper. University of Texas at Austin.

Lou, X. and H. Zhong (2003). When to Use Buy-price Auction?

Lucking-Reiley, D., D. Bryan, N. Prasad, and D. Reeves (1999). Pennies from eBay: the Determinants of Price in Online Auctions. Working paper. University of Arizona.

Lucking-Reiley, D. (2000a). Auctions on the Internet: What's Being Auctioned, and How? *Journal of Industrial Economics* 48. 227-252.

Lucking-Reiley, D. (2000b). Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-Commerce. *Journal of Economic Perspectives* 14(3). 183-192.

Maskin, E. (2004). The Unity of Auction Theory: Milgrom's Masterclass. *Journal of Economic Literature* XVII. 1102-1115.

Maskin, E. and J.G. Riley (1984). Optimal Multi-unit Auctions. In: F. Hahn (ed.). *The Economics of Missing Markets, Information and Games*. Oxford: Oxford University Press. 312-335.

Matthews, S. (1995). Comparing Auctions for Risk Averse Buyers: A Buyer's Point of View. *Econometrica* 55. 633-646.

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

Mathews, T. (2004). The Impact of Discounting on an Auction with a Buyout Option: a Theoretical Analysis Motivated by eBay's Buy-It-Now Feature. *Journal of Economics* 81(1). 25-52.

Mathews, T and B. Katzman. (2004). The Role of Varying Risk Attitudes in an Auction with a Buyout Option. mimeo. California State University-Northridge.

Matthews, T. (2005). Bidder Welfare in an Auction with a Buyout Option. http://buslab5.csun.edu/tmathews/pdf/BidderWelfare.pdf.

Mathews, T. and B. Katzman (2006). The Role of Varying Risk Attitudes in an Auction with a Buyout Option. *Economic Theory* 27. 597-613.

McAfee, R.P. (1993). Mechanism design by competing sellers. *Econometrica* 61. 1281-1312.

McAfee, R.P., and J. McMillan (1987). Auctions with Entry. Economics Letters 23 (4), 343-347.

McAfee, R.P. and J. McMillan (1996). Analyzing the Airwaves Auction. *Journal of Economic Perspectives* 10(1). 159-175.

McAfee, R.P., D. Quan, and D. Vincent (2000). How to Set Minimum Acceptable Bids, With an Application on Real-Estate Auctions. Mimeo

Melnik, M. and J. Alm (2002). Does a seller's e-commerce reputation matter? – Evidence from eBay auctions. *The journal of industrial economics* 50(3). 337-349.

Menezes, F.M. and P.K. Monteiro (2005). An Introduction to Auction Theory. Oxford University Press.

Milgrom, P. (2004). Putting Auction Theory to Work. Cambridge University Press.

Milgrom, P.R. and R.R. Weber. (1982). A Theory of Auctions and Competitive Bidding. *Econometrica* 50(5). 1089-1122.

Milgrom, P.R. (forthcoming). Package Auctions and Package Exchanges: the 2004 Fisher-Schultz Lecture. *Econometrica*.

Morgan, J., K. Steiglitz and G. Reis (2003). The spite motive and equilibrium behavior in auctions. *Contributions to Economic Analysis & Policy* 2(1), Article 5. http://www.bepress.com/bejeap/contributions/vol2/iss1/art5).

Myerson, R.B. (1981). Optimal Auction Design. Mathematics of Operations Research 6. 58-73.

Nagareda, T. (2003). Announced Reserve Prices, Secret Reserve Prices, and Winner's Curse. http://www.e.u-tokyo.ac.jp/cirje/research/workshops/micro/micropaper02/gs/nagareda.pdf

Namazi, A. (2005). Emergent Behavior and Criticality in Online Auctions. Ph.D. Dissertation. University of Cologne.

National Consumer League (2002). http://www.nclnet.org/shoppingonline/auctionsurvey.htm

Ockenfels, A. (2003). Reputationsmechanismen auf Internet-Marktplattformen: Theorie und Empirie. *Zeitschrift für Betriebswirtschaft*73(3). 295-315.

Ockenfels, A. (2003). New Institutional Structures on the Internet: The Economic Design of Online Auctions. In: M.J. Holler, H. Kliemt, D. Schmidtchen, and M. Streit (eds.). *Jahrbuch für Neue Politische Ökonomie* 20. 57-78. Tübingen: Mohr Siebeck.

Ockenfels, A. (2005). Bidding in Private Value eBay Auctions: A Controlled Field Experiment. Mimeo. Economics Department. University of Cologne.

Ockenfels, A. and A. E. Roth (2002). The Timing of Bids in Internet Auctions: Market Design, Bidder Behavior, and Artificial Agents. *Artificial Intelligence Magazine* Fall 2002. 79-87.

Ockenfels, A. and A. E. Roth (forthcoming). Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction. *Games and Economic Behavior*.

Ockenfels, A. and R. Selten (2005). Impulse Balance Equilibrium and Feedback in First Price Auctions. *Games and Economic Behavior* 51. 155-179.

Ow, T. T., C. A. Wood (2004). Trust Versus The Winner´s Curse – The Effects of Bidder Experience on Prices in Online Auctions. Working Paper. University of Notre Dame. http://www.nd.edu/~cwood1/research/BidderX.pdf.

Park, Y.-H., and E. Bradlow (2004). An Integrated Model for Whether, Who, When and How Much in Internet Auctions. Working paper. Wharton School, University of Pennsylvania.

Peters, M. and S. Severinov (1997): Competition Among Sellers who offer Auctions Instead of Prices. *Journal of Economic Theory* 75(1). 141-179.

Peters, M. and S. Severinov (forthcoming): Internet Auctions with Many Traders. *Journal of Economic Theory*.

Rasmusen, E. (2003). Strategic Implications of Uncertainty Over One's Own Private Value in Auctions. Working paper. Indiana University. http://php.indiana.edu/~erasmuse/papers/auction.pdf.

Reiley, D. H. (forthcoming). Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet. *RAND Journal of Economic Research*.

Reiss, J.P. (2004). Optimal auctions with public outside options. Working paper. University of Magdeburg.

Reiss, J.P. and J.R. Schoendube (2002). Entry and bidding behavior in sequential procurement auctions. Working Paper. University of Magdeburg.

Resnick, P. and R. Zeckhauser (2002). Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. In: M. R. Baye (editor). *The Economics of the Internet and E-Commerce*, Advances in Applied Microeconomics 11. Amsterdam. Elsevier Science.

Reynolds, S. S., J. Wooders (2003). Auctions with a Buy Price. Working Paper. University of Arizona. http://www.u.arizona.edu/~sreynold/buynow.pdf.

Reynolds, S. S. and John C. Wooders. Auctions with a Buy Price (2004). Eller College Working Paper No. 1010-05. http://ssrn.com/abstract=655362.

Riley, J.G. and W.F. Samuelson (1981). Optimal Auctions. *American Economic Review* 71. 381-392.

Roth, A. E. (2002). The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics. *Econometrica* 70(4). 1341-1378.

Roth, A. E. and Ockenfels, A. (2000). Last Minute Bidding and the Rules for Ending Second-Price Auctions: Theory and Evidence from a Natural Experiment on the Internet. NBER working paper 7729.

Roth, A. E. and A. Ockenfels (2002). Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet. *American Economic Review* 92(4). 1093-1103.

Roth, A. E., J.K. Murnighan, and F. Schoumaker (1988). The Deadline Effect in Bargaining: Some Experimental Evidence. *American Economic Review* 78(4). 806-823.

Rothkopf, M.H., T.J. Teisberg, and E.P. Kahn (1990). Why are Vickrey Auctions Rare? *Journal of Political Economy* 98(1). 94-109.

Samuelson, W. F. (1985). Competitive Bidding with Entry Costs. *Economics Letters* 17, 53-57.

Schindler, J. (2003). Late Bidding on the Internet. Working paper. University of Vienna.

Seifert, S. (2006). Posted Price Offers in Internet Auction Markets. Lecture Notes in Economics and Mathematical Systems. Springer Verlag, Berlin.

Shah, H. S., N. R. Joshi, and P. R. Wurman (2002). Mining for Bidding Strategies on eBay. Working Paper. Department of Computer Science. North Carolina University.

Shmueli, G., R. P. Russo, and W. Jank (2005). Modeling Bid Arrivals in Online Auctions. Working paper. Robert H. Smith School of Business. University of Maryland.

Simonsohn, U. (2005). Excess Entry Into High-Demand Markets: Evidence from the Timing of Online Auctions. Working paper. The Wharton School. University of Pennsylvania.

Simonsohn, U. and D. Ariely (2004). e-Bay's Happy Hour: Nonrational Herding in On-Line Auctions. Working Paper. The Wharton School. University of Pennsylvania.

Sinha, A. R. and Eric A.Greenleaf (2000). The Impact of Discrete Bidding and Bidder Aggressiveness on Sellers' Strategies in Open English Auctions: Reserves and Covert Shilling. *Marketing Science* 19(3). 244-265.

Smith, V.L. (1976). Experimental economics: Induced value theory. *American Economic Review*, Papers and Proceedings. 274-279.

Standifird, S., M. Roelofs, and Y. Durham (2004) eBay's "Buy-It-Now" function and the impact on bidder behavior. International Journal of Electronic Commerce 9(2). 167-176

Stryszowska, M. (2005a). Late and Multiple Bidding in Simultaneous and Overlapping Second Price Internet Auctions. CentER discussion paper.

Stryszowska, M. (2005b). On the Ending Rule in Sequential Internet Auctions. CentER discussion paper.

Stryszowska, M. (2005c). Coordination Failure in Internet Auctions with Costly Bidding. CentER discussion paper.

Tucker, J. M. and V. J. Massad (2003). An Examination of Seller Pricing Options at Online Auctions. Working Paper #2003-01. Shippensburg University. http://grove.ship.edu/research/wp/03tucker.pdf.

Tung Y.A., R.D. Gopal, A.B. Whinston (2003). Multiple Online Auctions. *IEEE Computer*. 88-90.

Varian, Hal R. (2006). Position Auctions. Working paper. Berkelee.

Vickrey, W. (1961). Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance* 16. 8-37.

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

Wan, W., T. Hock-Hai, and L. Zhu (2003). Determinants of Buy-It-Now Auction Choice on eBay. Conference Paper. 7th Pacific Asia Conference on Information Systems, Adelaide, Australia. http://www.pacis2003.org/proceedings/authors.html#D.

Wang, T. J. (2003). Is Last Minute Bidding Bad? Working paper. UCLA.

Wang, W., Z. Hidvégi, and A. B. Whinston (2001). Shill Bidding in English Auctions.

Wang, W., Z, Hidvégi, and Andrew B. Whinston (2004). Shill-Proof Fee (SPF) Schedule: the Sunscreen against Seller Self-Collusion in Online English Auctions. http://www.wenli.net/Application/Shill.pdf.

Wang, X., A. L. Montgomery, and K. Srinivasan (2004). When Auction meets Fixed Price: A Theoretical and Empirical Examination of Buy-it-Now Auctions. Working Paper. Carnegie Mellon University. http://www.andrew.cmu.edu/user/alm3/papers/buyitnow.pdf

Ward, S. G. and J. M. Clark (2002). Bidding Behavior in On-line Auctions: An Examination of the eBay Pokemon Card Market. *International Journal of Electronic Commerce* 6(4). 139-155.

Wilcox, R. T. (2000). Experts and Amateurs: The Role of Experience in Internet Auctions. *Marketing Letters* 11(4). 363-374.

Wintr, L. (2004). Some Evidence on Late Bidding in eBay Auctions. Working paper. Clark University, Worcester MA.

Wolf, J. R., H. R. Arkes, and W. A. Muhanna (2005). Is Overbidding in Online Auctions the Result of a Pseudo-Endowment Effect? Working paper. Ohio State University.

Wolfram, C. D. (1998). Strategic Bidding in a Multiunit Auction: An Empirical Analysis of Bids to Supply Electricity in England and Wales. *RAND Journal of Economics* 29. 703-725.

Wooders, J. and S. S. Reynolds (2004). Auctions with a Buy Price. *Econometric Society 2004 North American Summer Meetings 130, Econometric Society.*

Yin, P.-L. (2005). Information Dispersion and Auction Prices. Working Paper. Stanford University.

Zeithammer, R. (2003). Forward-looking bidders in online auctions: A theory of bargain-hunting on eBay. Working Paper. University of Chicago.

Zheng, C. (2002). Optimal Auction with Resale. *Econometrica* 70 (6), 2197-2224.