Sequencing and Information Revelation in Auctions for Imperfect Substitutes: Understanding eBay's Market Design

Eric Budish[1]

September 12, 2008

[1]Contact: Harvard University; Economics Department; Littauer Center; Cambridge, MA 02138. ebudish@hbs.edu. I thank Susan Athey, Lucas Coffman, Itay Fainmesser, Drew Fudenberg, Ian Jewitt, Adam Juda, Robin Lee, John Lewis, Paul Milgrom, Jerry Orloff, Marco Pagnozzi, Ariel Pakes, David Parkes, Robert Ritz, Laura Serban, and especially Paul Klemperer and Al Roth for helpful discussions.

Abstract

An auction marketplace like eBay consists of (i) an individual auction design; and (ii) a multi-auction platform design, i.e., the set of rules that organize and convey information about multiple individual auctions. This paper shows that two aspects of eBay's platform design – sequencing of auctions by unique ending time, and provision of information about both current and near-future objects for sale – substantially increase the social surplus generated by single-unit second-price auctions when the goods traded are imperfect substitutes. The remaining inefficiency from not using a multi-object auction is surprisingly small.

JEL Classification: D44, D82, D83, L13, L81

Keywords: Market Design, Auctions, Information in Auctions, Entry in Auctions, Sequential Auctions, eBay, e-commerce, Platforms

The internet auctioneer eBay was created in 1995 with collectibles traders in mind: its auction format was a natural adaptation of that used by art auctioneers such as Sotheby's and Christie's, and the company claimed that the founder's wife's desire to trade collectible "Pez Dispensers" inspired its creation. The character of the goods traded on eBay has changed dramatically. Now, most of the items auctioned are consumer goods that have close substitutes; used automobiles alone constitute around one-quarter of the approximately $60bn of annual trading volume on eBay.[1]

Modern auction theory suggests that eBay (or an entrant) should consider adopting a multi-object auction format – e.g., the simultaneous ascending auction (SAA) or Vickrey-Clarke-Groves mechanism (VCG) – for the trade of goods with close but imperfect substitutes. Of course, while these mechanisms are efficient (i.e., social-surplus maximizing), they also may be complex for participants to understand and for auctioneers to implement. In high-stakes contexts such as spectrum allocation it may be reasonable to ignore these particular efficiency-complexity tradeoffs; e.g., bidders can hire sophisticated auction consultants for a presumably negligible fraction of the value of the goods being auctioned. But in low-stakes contexts such as eBay it is worth understanding whether simpler mechanisms perform satisfactorily.

To evaluate this question we need to analyze how single-unit auctions aggregate up into a multi-unit auction marketplace. In addition to its single-unit auction design, eBay adapted two aspects of Sotheby's and Christie's *multi-auction platform design*: individual auctions are sequenced (by unique ending time), and bidders are allowed to evaluate all of the objects in the sequence before deciding whether and how much to bid in any particular auction.

This paper shows that this combination of single-unit auction design and multi-auction platform design may approximate the efficiency performance of a more sophisticated auction design. (The purpose is *not* to explain eBay's success; surely, sensible market design is just one component of the story).

Both sequencing auctions and revealing information seem like natural design decisions, perhaps in part because they have been used by the classic auction houses. Both features are also present in the market design used by the largest wholesale used-auto auctioneer,

[1] The Pez Dispenser story turned out to be apocryphal. See Cohen (2002).

eBay (2006a) provides trading volume by category for the year 2005. (eBay appears to have subsequently stopped providing this level of category volume detail). The categories with in excess of $2bn are: Motors ($13.6bn), Consumer Electronics ($3.5), Clothing & Accessories ($3.4), Computers ($3.1), Books/Music/Movies ($2.6), Home & Garden ($2.5), Collectibles ($2.2), and Sports ($2.1).

Trading volume by category is not available for the early days of eBay. However, the following excerpt from the introductory section of its 1998 IPO prospectus is illustrative of its focus on collectible goods:

"eBay ... [is] ... a Web-based community in which buyers and sellers are brought together in an efficient and entertaining auction format to buy and sell personal items such as antiques, coins, collectibles, computers, memorabilia, stamps, and toys." (eBay, 1998).

Manheim (see Lewis, 2008). But there are alternative auction marketplaces that have made quite different decisions.

For instance, in 2006 the online event-ticket marketplace StubHub ran the first ever auction for all of the tickets to a single event. Its single-unit auction design was similar to eBay's (a unit is a pair of tickets). But it aggregated these auctions into a multi-auction platform in an importantly different way: the auctions had identical hard-close ending times (see Harrington, 2006; Smith, 2006). Since most of the bidding in online auctions with hard closes occurs in the final moments (see Roth and Ockenfels, 2002), buyers could effectively participate in just a single of the auctions. Rather than *Sequencing* the auctions, StubHub effectively *Separated* them. Auctions at charity benefits are often organized similarly, which may explain why the final moments of so-called "silent auctions" are often anything but.

Engelbrecht-Wiggans (1994) describes a sequential auction for restaurant equipment in which future items are hidden behind a curtain until their turn for sale. Some gimmicky online auction sites with very short auction durations (e.g. Bidz.com[2]) also suppress information about future objects for sale. Rather than *Revealing* information about future objects, these platform designs *Hide* such information.

These alternative designs motivate the study of a simple two-by-two taxonomy of multi-auction platforms. The individual auctions are either Sequential or Separated; information about future objects is either Revealed or Hidden. The Separated condition can be interpreted more broadly as a tractable device for modeling an auction platform in which there are obstacles to participating in multiple individual auctions.[3] The Hidden condition can be interpreted more broadly as a modeling convention that captures that auction buyers know they have an outside option but are not yet sure of its value.

The main theoretical result of this paper (Theorem 1) is that the Sequential Auction with Information Revealed is the most efficient in the taxonomy. This result obtains for any number of unit-demand bidders and any (continuously differentiable) distribution of bidders' values. A limitation of the theoretical analysis is that we restrict attention to platforms that aggregate exactly two individual auctions, for tractability.

We then use numerical analysis to give a sense of magnitudes. The Sequential Auction with Information Revealed generates at least 99% of the efficient social surplus over all sim-

[2]On Bidz.com "auctions start at $1 every 5 seconds!" $132 million of jewelry was sold on Bidz.com in 2006.

This kind of information suppression occurs in the context of non-auction platforms as well; see Hagiu and Jullien (2008).

[3]For instance, both Craigslist and Google Base are frequently used for auctions, but listings are sequenced by starting time and search closeness-of-fit, respectively. So, a buyer might participate in one auction, lose, and then learn that some other item he is interested in has meanwhile been sold. The soft closing times of the now-defunct Amazon and Yahoo! auction marketplaces might have the same effect of making it more difficult to participate in the complete sequence of auctions.

ulated value distributions. By contrast, the other formats can have inefficiency that is an order of magnitude worse. The insight that explains the consistently strong efficiency performance of the Sequential Auction with Information Revealed is that the performance of the alternative mechanisms is in a sense negatively correlated (Theorem 1 then tells us that the eBay design performs better than the best alternative in the taxonomy). Revealing information is especially important for welfare when the distribution of bidder values is heavily left skewed (high values are rare), whereas Sequencing auctions is especially important when the distribution is heavily right skewed (high values are common).

Numerical analysis also allows us to provide a simple robustness check to ensure that the results do not depend on the assumption that there are exactly two objects. The results appear to be robust to this extension, both qualitatively and quantitatively.

This paper emphasizes efficiency because an auction marketplace that is sufficiently inefficient (e.g., taking switching costs into account) should be vulnerable to entry: a competitor could offer both sides of the market strictly more surplus (see Ellison, Fudenberg, and Mobius, 2004 for a study of competition between auction marketplaces). This emphasis is standard for auction markets with multiple buyers and multiple sellers (e.g., assignment, double auctions). However, it may be important to consider revenues as well.[4] For most simulated value distributions Sequencing and Revealing Information increase revenues, and it can be shown theoretically that revealing information in a sequential auction increases all types of bidders' first-round bids and increases the strength of the second-round bidder pool (competition effects).

A metric that is related to revenues, and that eBay emphasizes in its financial reporting and marketing materials, is the "seller success rate". In the model of this paper, in which sellers' costs are strictly lower than the base of the bidder value support and strategic reserve prices are not allowed, the Sequential Auction with Information Revealed achieves a perfect success rate of one, like an efficient multi-object auction and unlike all other members of the taxonomy.

Together these results suggest that it is not surprising per se that a non-combinatorial auction marketplace like eBay's has been successful at facilitating the trade of goods with imperfect substitutes. A simple single-object auction format, combined with a smart platform design that organizes multiple of these single-object auctions in an economically useful way, performs nearly as a more complicated alternative.

These results also suggest that it is not surprising that StubHub's ticket auction was

[4] Bulow and Klemperer (2007) emphasize that in the standard single-seller auction context – e.g., the one-time sale of a valuable asset, like a company – revenue is probably the more important metric, because revenue is what will incentivize sellers to create valuable assets in the first place.

unsuccessful, and that StubHub appears to have abandoned the format.[5] Nor might it be surprising that Bidz.com has allegedly had to use shill bidding to increase its success rate. (Miniter, 2008)

In addition to highlighting a partial explanation for eBay's success this paper also provides a framework for understanding an important vulnerability. Some recent empirical studies of eBay auctions suggest that many bidders for substitute goods bid in just a single auction. For instance, Juda (2005) finds that over half of the bidders in his study (of auctions for a particular model of computer monitor) participate in exactly one auction, are unsuccessful, and then exit the market. Lee and Malmendier (2006), too, document a failure to participate in multiple auctions, and Bajari and Hortacsu (2003) find that bidder participation costs dampen entry.

In the language of this paper's taxonomy these empirical findings suggest that eBay auctions may be more Separate than Sequential. There may indeed be ways to meaningfully enhance the efficiency of eBay's auctions for substitutes, but this paper suggests we look for them not in multi-object auctions per se, but in technologies that enable bidders more easily to express substitutes preferences over multiple objects. This appears to be eBay's understanding as well, as evidenced by its recent enhancements to its bidding proxy to enable bidders to participate in the full sequence of auctions with just a single visit to their computer.[6]

Related Literature This paper represents a first foray into the study of multi-auction platform design: taking the individual auction design as given, we study the rules that organize and convey information about multiple of these individual auctions. Methodologically, it is written in the spirit of the emerging literature on market design that looks at the institutional details of markets, complementing theory with simulation, experimental or case-study evidence to give a sense of magnitudes. (Klemperer, 2002; Roth, 2002)

This paper is most closely related to the literatures on combinatorial auctions (e.g., Milgrom, 2004; Cramton et al, 2006) and eBay (e.g., Bajari and Hortacsu 2003, 2004; Lewis, 2006). The literature on eBay as a market institution has focused mainly on the analysis of individual auctions, whereas the combinatorial auctions literature has focused on the design of auction mechanisms to handle allocation problems considerably more complex and

[5] The StubHub auction's average selling price was $50 per ticket, versus a $148 average aftermarket value for tickets for that particular tour. Some of the StubHub auctions closed at as low as $3 per ticket. Nevertheless StubHub's CEO described the auction as a "successful experiment in true dynamic pricing." See Cohen and Grossweiner (2006).

eBay acquired StubHub in early 2007 and currently operates it as a standalone business.

[6] This new feature is called "Bid Assistant", which eBay indicates will save "busy buyers" "valuable time". See http://pages.ebay.com/bidassistant.

high-stakes than the unit-demand environment of interest here. This paper articulates a middle ground between these literatures, by examining how individual single-unit auctions can aggregate up into a multi-object auction marketplace. For the case of interest, such a marketplace may approximate the benefits of a more-sophisticated combinatorial auction, and may be easier to understand by low-stakes buyers and sellers.

This paper also is related to the auction-theoretic literatures on entry and information disclosure. In this paper, sequencing is of first-order importance for a fixed set of entrants; previous results have indicated that entry is of first-order importance when the set of participating bidders is endogenous (Athey et al, 2004; Klemperer, 2002; Bulow and Klemperer, 1996). In this paper, providing buyers with additional private information benefits the sellers (usually) and social surplus (always). This is distinct from the Linkage Principle (Milgrom and Weber, 1982) which indicates that providing additional *public* information benefits sellers without affecting social surplus, though it has similar implications for auction policy.

Finally, this paper is related to the burgeoning literature that studies the design of two-sided platform markets (for a recent survey see Rochet and Tirole, 2006). The two-sided platform literature has generally focused on how platforms should set prices (or subsidies) on each side of the market, taking into account that participants on each side generate a network externality for those on the other. This literature typically abstracts from the details on how the interaction between the two sides creates value (notable recent exceptions are Athey and Ellison, 2008, and Hagiu and Jullien, 2008). By contrast, this paper takes the set of participants as exogenous and focuses on how the design of the platform affects how much value the participants are able to create. Work that combines both approaches is likely to be a fruitful avenue for future research.

Organization of the Paper The remainder of this paper is organized as follows. Section 1 presents the model. Section 2 considers equilibrium bidding behavior for each of the auction formats in the two-by-two taxonomy, drawing on Zeithammer (2006) for the analysis of the sequential auction with information revealed.[7] Section 3 presents the main efficiency results, making novel use of a classic result from the theory of auctions with entry, due to McAfee and McMillan (1987) and Engelbrecht-Wiggans (1993). Section 4 presents revenue results. Section 5 presents numerical simulations for various parameterizations of bidders' values, to give a sense of magnitudes. Section 6 concludes. Proofs and a simple two-value example are

[7] I became aware of Zeithammer (2006) after developing the main results of this paper for a two-value model, for which bid functions can be obtained in closed form. Zeithammer provides an equilibrium existence theorem for Sequential Auctions with Information Revealed (stated herein as Lemma 1) that enabled a generalization of this paper's analysis of efficiency and revenue differences across market designs.

There is no overlap in our papers' motivations or results. Essentially, the papers are complements, not substitutes.

contained in appendices.

1 The Model

Basic Setup and Design Taxonomy. Two items, $j = 1, 2$, are sold by second-price sealed-bid auction to a group of $n \geq 3$ bidders. The sealed-bid assumption simplifies the analysis and is motivated by the fact that most bidding on eBay occurs towards the very end of the auction (Roth and Ockenfels, 2002). The restriction to two items is purely for tractability, though it is relaxed in numerical analysis (Section 5.2.5).

The auctions are either *sequential*, which has the usual meaning, or *separate*, which means that each bidder must choose whether to participate in the auction for 1 or the auction for 2. The entry decisions in the separate auctions are simultaneous, meaning that bidders cannot observe others' entry choices before making their own.

We are interested in two informational environments, called *revealed* and *hidden*. If information is revealed, each bidder i learns his private value for each object before he makes any bidding or entry decisions. When information is kept hidden, each bidder i learns his private value for object j only at the time of its auction. Specifically: in the sequential auction, bidder i learns only his first-round value before he must make his first-round bid; in the separate auctions, bidder i learns his value for the one auction he enters only after he has made his entry decision.

Buyers. Bidders are risk-neutral and have unit demand; in particular, the winner of the first of two sequential auctions is assumed to exit the game (this is like assuming high-enough resale costs). Bidders who lose both auctions have an outside option normalized to zero. Bidder i's private value for object j, v_{ij}, is independently and identically distributed according to the probability distribution $f^j(v)$, with support on $[\underline{v}, \overline{v}]$, $\underline{v} < \overline{v}$. We assume that the objects are stochastically equivalent, i.e., that $f^1(v) = f^2(v) \equiv f(v)$ and that $f(\cdot)$ is a continuously differentiable density with full support on $[\underline{v}, \overline{v}]$. It is possible to add a constant term ξ_j (e.g. representing object j's quality) to all bidders' values for a particular object j; the only effect this will have in equilibrium is to increase all bids for object j by exactly ξ_j without any affect on the allocation or the bidding for the other object.

The key assumption about buyers is that bidder i's value for object j is independent of all other bidders' idiosyncratic preferences (including his own for the other object). Independence across bidders is required to make use of Zeithammer's (2006) equilibrium existence result for the sequential auction with information revealed. Independence within bidders across objects allows the comparison across the two information environments to be the most meaningful. For instance, if the bidders' values are perfectly correlated across objects then the hidden and revealed environments will be outcome equivalent.

Sellers. Sellers play a passive role. All they do is set a minimum bid equal to their salvage value of $c < \underline{v}$. In particular, they do not use strategic reserve prices or buy prices.

This assumption is reasonable in cases where the seller has imprecise information about the distribution of buyers' values. Theory suggests that sellers who have precise information about buyers' values will use buy prices in most natural low-stakes auction markets.[8] Indeed, they are used in about one-third of eBay auctions.[9]

Model Applicability. The model applies best to markets in which goods have some close substitutes, but that are not so thick that the pricing problem becomes trivial. Markets for which the model appears a particularly good approximation are used, discontinued, or supply-constrained goods (e.g., automobiles, business equipment, premium tickets to sold-out events). The model is not intended to apply to markets for new goods without supply constraints (e.g., computer parts). In such cases, sellers will typically use fixed prices and may have access to additional supply, and it seems unrealistic to assume that buyers' outside options are symmetric.

2 Equilibrium Bidding

2.1 Sequential Auctions

Suppose that bidders submit their first-round bids for the sequential auction with information revealed symmetrically according to a function $b_1^r(v_1, v_2)$, and suppose that the highest bid out of a particular set of n bidders is c_1. Bidder i's continuation surplus, if he is one of the $n-1$ losers, is the expected profit of a bidder whose value is v_{i2} at the dominant strategy equilibrium of a single second-price auction, when his $n-2$ opponents' types are random draws from the set $\{(v_1, v_2) : b_1^r(v_1, v_2) \leq c_1\}$ based on the underlying joint density $f(v_1, v_2)$.[10]

[8] A "buy price" is a price which, if bid, ends the auction immediately. Milgrom (2004) shows that a buy price increases expected revenues when there are moderate entry costs. Budish and Takeyama (2001) and Hidvegi et al (2006) show that a buy price increases expected revenues when buyers are risk averse. (In all three of these papers it is assumed that the seller knows the distribution of buyers' values.)

Intuition suggests that one of these conditions will typically be satisfied on eBay: when the stakes are low, entry costs are likely to matter; when the stakes are high, risk aversion will matter.

eBay's Buy It Now differs slightly from a traditional buy price, in that the option to transact at the BIN price is withdrawn upon the first bid greater than the seller's reserve.

[9] eBay (2006b) indicates that Buy It Now auctions account for 34% of trade by volume. Meeker (2006) indicates that BIN is used in 35% of auction listings, and tends to appear in a greater percentage of listings in conventional goods categories than in traditional collectibles categories.

[10] The conditional density $f(v_1, v_2 | c_1; b_1^r)$ is simply $\frac{f(v_1)f(v_2)}{\int_{b_1^r(x_1,x_2)\leq c_1} f(x_1)f(x_2)dx_1dx_2}$ for $b_1^r(v_1, v_2) \leq c_1$ and zero otherwise.

We can write this surplus as $S(v_{i2}, c_1; b_1^r)$. If b_1^r is the unique symmetric equilibrium b_1^{r*} (discussed in Section 2.1.2), we simply write $S(v_{i2}, c_1)$.[11]

As we will see, a single continuation surplus function $S(v_{i2}, c_1)$ will enable us to characterize bidding in the sequential auction for both information environments.

2.1.1 Equilibrium Bidding with Information Hidden

Suppose that bidders submit their first-round bids for the sequential auction with information hidden according to a function $b_1^h(v_1, v_2) = b_1^h(v_1, -)$ (with "$-$" indicating a value that is not yet known), and so the continuation surplus function is $S(v_{i2}, c_1; b_1^h)$. Since bidders' values are independent across rounds and bidders do not know their second-round values when they submit their first-round bids, it follows that c_1 is not informative about the continuing bidders' second-round values (no matter what b_1^h is played). That is, the conditional marginal distribution $f_2(v_1, v_2|c_1; b_1^h)$ is equal to the underlying bidder value distribution $f(v_2)$.

Observe that b_1^{r*} must be bounded above by $\overline{v}$; in equilibrium no bidder will submit a first-round bid in excess of his value. So the conditional marginal distribution $f_2(v_1, v_2|\overline{v}; b_1^{r*})$ must also be equal to the underlying value distribution $f(v_2)$: receiving information that all of one's second-round opponents bid weakly less than the maximum possible value in round 1 is equivalent to receiving no information at all about these same opponents. This implies that $S(v_{i2}, c_1; b_1^h) = S(v_{i2}, \overline{v}; b_1^{r*}) \equiv S(v_{i2}, \overline{v})$ for any b_1^h, and so the (unique Perfect Bayes-Nash) equilibrium bid function if information is hidden is:

$$b_1^{h*}(v_1, -) = v_1 - \mathbb{E}_{v_2} S(v_2; \overline{v}) \tag{1}$$

At this bid amount, rational bidders are indifferent between winning and losing at the margin.[12]

$S(v_2; \overline{v})$ can be calculated explicitly as $\frac{1}{n-1}\mathbb{E}(F_{(1:n-1)}(v) - F_{(2:n-1)}(v))$ - the expected difference between the highest and second-highest of $n-1$ bidder values, multiplied by the probability of being the bidder whose value is highest. If $\underline{v} - \mathbb{E}_{v_2} S(v_2; \overline{v})$ is strictly less than the seller's salvage value c, then some bidders will abstain from the first auction.

[11] Zeithammer (2006) contains explicit expressions for $S(v_2, c_1)$ in terms of the underlying joint density $f(v_1, v_2)$.

[12] Perfectness eliminates any asymmetric equilibria here. Participation in the second-round auction has the same expected value to all bidders, independent of the play of the first-round auction. So, the first-round auction is like a one-shot second-price sealed-bid auction in which losers receive a set payment.

2.1.2 Equilibrium Bidding with Information Revealed

When information is revealed different types will shade their bids by different amounts; all bidders bid so as to be indifferent between winning and losing at the margin. Equilibrium bidding, in the first of two second-price auctions with information revealed, is characterized by the following equation:

$$(2) \qquad \underbrace{b_1^{r*}(v_1, v_2)}_{\text{Bid}} = \underbrace{v_1}_{1^{\text{st}} \text{ pd value}} - \underbrace{S(v_2, b_1^{r*}(v_1, v_2))}_{\text{Surplus as a marginal loser}}$$

The key to equation (2) is that the bidder's own equilibrium bid is the argument in the $S(\cdot)$ function defining which types of first-round bidders continue into the second-round. A necessary condition for a bid to be a component of an equilibrium strategy is that it satisfy (2). The logic is a simple generalization of Vickrey's (1961). Given a highest outside bid x, i's bid b_{i1}^r affects neither the price he pays if he wins (x) nor his continuation value if he loses (he faces a set of opponents who each bid less than x in the first round). If $b_{i1}^r > v_{i1} - S(v_{i2}, b_{i1}^r)$ then i prefers to marginally lose rather than marginally win. Such a b_{i1} cannot be an equilibrium bid since i can improve his surplus by bidding a small amount less, i.e., by breaking ties in the way he prefers. A similar argument indicates that $b_{i1}^r < v_{i1} - S(v_{i2}, b_{i1}^r)$ cannot be an equilibrium.

Zeithammer (2006) provides sufficient conditions for the existence of a unique symmetric equilibrium bidding strategy b_1^{r*}.

Lemma 1 *(Zeithammer 2006; Proposition 3). If $f(\cdot)$ is a continuously differentiable density with full support on $[\underline{v}, \overline{v}]$, then there exists a unique symmetric pure-strategy Perfect Bayes-Nash equilibrium in which types with $v_1 > \underline{v}$ bid according to a unique $b_1^{r*}(v_1, v_2)$, and the zero-measure set of types with $v_1 = \underline{v}$ abstain from the first auction. The continuation-surplus function $S(v_2, c_1)$ is unique, and continuous in c_1 with $\frac{dS(v_2,c_1)}{dc_1} > -1$.*

2.2 Separate Auctions

Strategies in the Separate auctions consist of an entry decision followed by a bidding decision. Since bidders enter a single second-price sealed-bid auction the bidding decision is trivial in any Perfect equilibrium: bidders simply bid their values.

The entry decision is a bit more subtle, because bidders may want to coordinate. If we restrict attention to symmetric Perfect Bayes-Nash equilibria, then if information is revealed the unique[13] equilibrium entry decision of each bidder i is to enter auction j iff $v_{ij} > v_{i(3-j)}$,

[13] Bidders can break ties, which occur with probability zero, however they like. Otherwise, the equilibrium is unique.

and if information is hidden the unique equilibrium entry decision of each bidder i is to randomize uniformly between the two auctions.

There also exist asymmetric PBNEs which address the coordination problem, for both informational environments. For instance, if $n = 4$ and information is hidden, then it is an equilibrium for bidders 1 and 2 always to enter 1 and bidders 3 and 4 always to enter 2.

3 Efficiency

In the present context, in which payments are pure transfers and all agents are risk neutral, social surplus is defined as the sum of the winners' values for the objects they receive (if an object is unallocated, the seller, whose value is c, is the "winner"). A fully efficient auction marketplace is one that achieves the maximum possible expected social surplus, i.e., realizes the social-surplus maximizing allocation across all realizations of bidder values.

In order to make efficiency comparisons across auction formats we need to translate individuals' self-interested bidding decisions into aggregate social surplus contributions. The basic insight that enables this is a classic result from the theory of auctions with entry, due to McAfee and McMillan (1987) and Engelbrecht-Wiggans (1993). The latter statement is more general, and is summarized for the case relevant to the present analysis:

Lemma 2 *(Engelbrecht-Wiggans, 1993: Proposition 1) Let bidders' private values for a single object be drawn from a symmetric joint probability distribution. At the dominant strategy equilibrium of the second-price sealed bid auction, each bidder has an expected profit equal to his marginal contribution to social surplus.*

This result is useful here because it implies that the marginal loser surplus term $S(v_2, c_1)$ is equivalent to the expected contribution to second-round surplus of a bidder with value v_2, when the set of $n-2$ second-round opponents is drawn from those who bid weakly less than c_1 in the first auction in equilibrium.

3.1 Comparison Results

Our first comparison is across information environments for sequential auctions:

Proposition 1 *The unique symmetric PBNE of the Sequential Auction with Information Revealed generates greater expected social surplus than the unique PBNE of the Sequential Auction with Information Hidden.*

The idea of the proof is as follows. If information is hidden the allocation is "greedy" – the first object is allocated to whomever of the n bidders has the highest value for it, and then the second object is allocated to whomever of the $n-1$ remaining bidders has the highest value for it.

If information is revealed the allocation is non-greedy whenever the bidder (i) with the highest first-round value shades his bid sufficiently that he loses the first auction to some other bidder (j) who has a lower first-round value but shades less.

Sometimes i's losing the first auction turns out ex-post to be social-surplus maximizing, and sometimes not (e.g., if some third bidder k has a very high second-round value).

What Lemma 2 tells us is that i's and j's shading decisions in round 1 are directly related to their *expected* social-surplus contribution in round 2. The proof combines this insight with the bid functions (2) to create a surplus-crediting scheme that yields the desired result; essentially, given the equilibrium bidding functions, the fact that i loses round 1 despite having the highest value is more than compensated by the fact that he participates in round 2 instead of j.[14]

Our next comparison is across auction formats, for either information environment:

Proposition 2 *For either Hidden or Revealed information, the unique symmetric PBNE of the Sequential Auction generates greater expected social surplus than any PBNE of the Separate Auctions.*

There exist value realizations for either information environment in which Separate auctions generate greater social surplus than Sequential auctions. Nevertheless, the proof for the Hidden information case is trivial: expected surplus is the largest of n and $n-1$ draws in the Sequential auction, versus $n-k$ and k draws (for some random k) in the Separate auctions.

The proof for the Revealed information case is more subtle, and has a similar structure to that of Proposition 1. Instead of focusing on the outcomes of the bidder with the highest value for the first object, it focuses on the outcomes of the bidder with the highest bid for the first object in the sequential auction, again using (2), Lemma 2, and a careful surplus-crediting scheme.

Notice that Proposition 2 allows for any PBNE of the Separate Auctions, not just the unique symmetric PBNE. This is important because asymmetric equilibria seem particularly plausible for Separate Auctions, due to the presence of coordination problems. Note too that no analogue of Proposition 1 is stated for the Separate Auctions. It is of course true

[14] The equilibrium bidding strategy implicit in (2) requires that bidders understand the informational content of their opponents' bids. Eyster and Rabin (2005) argue that in various economic environments this is unrealistic. Reassuringly, if bidders are "naive" and bid according to $b_1^{naive} = v_1 - S(v_2, \overline{v})$ – that is, they ignore the informational content of marginally losing – then an appropriately restated Proposition 1 still obtains. The proof is analogous, and is omitted.

that if we limit attention to symmetric PBNEs, then revealing information increases the expected social surplus of the Separate Auctions.

Together Propositions 1 and 2 imply that the eBay multi-auction platform design is the most efficient of the taxonomy (though of course not fully efficient):

Theorem 1 *The unique symmetric PBNE of the Sequential Auction with Information Revealed generates greater expected social surplus than any PBNE of any of the alternate formats in the taxonomy.*

Remark 1 *The Sequential Auction with Information Revealed is not fully efficient. There exist multi-object auction formats (e.g., the Simultaneous Ascending Auction or Vickrey-Clarke-Groves mechanism) that are fully efficient in this environment.*

If $n = 3$ and values are distributed uniformly on $[0, 1]$ then the worst-case efficiency loss for the Sequential Auctions with Information Revealed occurs when the values are $\{(1, .99); (0, 1); (.58, 0)\}$. The bidder whose values are $(1, .99)$ just loses the first auction to the $(.58, 0)$ bidder, for 0.42 of squandered social surplus. The worst-case efficiency loss of each of the other formats in the taxonomy is 1. (This point is elaborated substantially in Section 5.2.1).

This inefficiency is in contrast to results in Weber (1983) and Peters and Severinov (2006) for homogeneous objects, and Beggs and Graddy (1997) for vertically differentiated objects. In these environments: 1) one-dimensional bids are sufficient statistics for multi-item demand profiles; and 2) social surplus is always maximized by the greedy allocation.

4 Revenues

There is no uniquely best market design in the taxonomy from the sellers' perspective, in contrast to the finding in Theorem 1 that the eBay design is the unique best from society's perspective. What follows are several results that are suggestive of the Sequential Auction with Information Revealed having the strongest revenue performance, as well as an examination of cases in which other designs outperform.

4.1 Revenue Consequences of Revealing Information

We begin with two results that indicate that second-round competition in the sequential auction with information revealed is "stronger than random". Proposition 3 takes the perspective of the second-round seller, whereas Lemma 3 takes the perspective of buyers anticipating their second-round surplus.

Proposition 3 *(Second-Round Competition Effect) If information is revealed in a Sequential Auction then the distribution of participating buyers' second-round values first-order stochastically dominates $f(v_2)$, the unconditional distribution. Consequently, expected second-round revenue is higher in the Sequential Auction with Information Revealed than in the Sequential Auction with Information Hidden.*

Lemma 3 *For any c_1, $v_2 : S(v_2, c_1) \leq S(v_2, \overline{v})$.*

The stochastic dominance in Proposition 3 follows immediately from $\frac{db_1}{dv_2} < 0$ and the independence of values across objects. Revenue is higher because first-order stochastic dominance of the participating bidder-value distribution implies first-order stochastic dominance of the second-highest of $n-1$ draws from the bidder-value distribution.

Our next result, which derives from Lemma 3, shows that first-round competition in the sequential auction with information revealed is strong as well.[15]

Proposition 4 *(First-Round Competition Effect) Revealing Information in a Sequential Auction increases each bidder's first-round bid in expectation. That is, for any v_{i1}, $\mathbb{E}_{v_{i2}}[b_1^{r*}(v_1, v_2)] \geq b_1^{h*}(v_1, -)$.*

Note that Proposition 4 is not a statement about first-round revenues, which depend on the second-highest bid. While adding positive-mean noise to a distribution always increases the expectation of the first (highest) order statistic, it need not increase the expectation of the second order statistic.

Again, we do not make a revenue comparison across information environments for the Separate Auctions due to the multiplicity of equilibria. If we limit attention to the unique symmetric PBNEs, then revealing information increases expected revenue in the Separate Auctions.

4.2 Revenue Consequences of Sequencing

Sequencing increases participation: bidders participate in $2 - \frac{1}{n}$ auctions on average, versus one for the Separate auctions. Proposition 2 shows that this increase in participation unambiguously benefits social surplus. It is not always the case, however, that Sequencing increases revenues. There is a tradeoff: Sequencing increases participation, but bidders in the first auction of the sequence shade their bids.

Whether participation or shading is more important depends on the order statistics of the bidder value distribution. Let $F(v)$ denote the cumulative distribution function of the

[15] Note that the naive bidders discussed in footnote 20 avoid the first-round competition effect of Proposition 4: Eyster and Rabin's "cursed" bidders are "blessed" in this environment.

bidder value distribution $f(v)$, and let $V \sim F(v)$ and $V^2 \sim [F(v)]^2$. V^2 is the cumulative distribution function for the highest of two independent draws from V. We write $\mathbb{E}(V_{k:n})$ and $\mathbb{E}(V^2_{k:n})$for the expectation of the k^{th} highest of n draws from V and V^2, respectively.

If information is revealed, the revenue benefit from increased participation is closely related to $\mathbb{E}(V_{2:n}) - \mathbb{E}(V^2_{2:(n/2))}) > 0$. This expression indicates the expected increase in the second-highest bidder value from adding back $n/2$ lower-of-2 draws to the $n/2$ higher-of-2 draws that constitute the bidding pool for separate with information revealed. If information is hidden, the relevant expression is instead $\mathbb{E}(V_{2:n}) - \mathbb{E}(V_{2:(n/2))}) > 0$, because the $n/2$ participants in each separate auction are random.[16]

The revenue costs from shading are more closely related to the expected difference between the first- and second-highest of n bidder values, $\mathbb{E}(V_{1:n}) - \mathbb{E}(V_{2:n}) < 0$. The amount by which bidders shade is their conditional expected profit in the second auction, which is their expectation of the difference between their value and the second-highest value, conditional on their value being the highest of the $n-1$ remaining bidders.

An example of a distribution for which sequencing harms total revenues for either information environment is the two-value distribution (defined formally in Appendix B) with $n = 4$, and p, the probability that a bidders' value is Low, equal to 0.9. The uniform distribution is an example for which sequencing increases revenues.

4.3 Seller Success Rates

A metric that eBay emphasizes as of central importance is the proportion of auctions that result in a successful sale, which eBay calls the "conversion rate".[17] Though it has not been possible to compare revenues across the taxonomy, it is possible to conclude that eBay's auction format has the highest conversion rate in the taxonomy.

Proposition 5 *For any n, c and any distribution of bidder values, the unique PBNE of Sequential Auctions with Information Revealed has a seller success rate of one. For each of the other auction formats in the taxonomy there exist parameters for which the success rate in symmetric PBNE is strictly less than one.*

The success rate of one for eBay auctions follows immediately from Proposition 1 of Zeithammer (2006). The success rate of the Sequential Auction with Information Hidden

[16] These expressions ignore the variance in the number of participants in each separate auction, and so may understate the benefit from participation. This variance plays a central role in Section 5.2.1.

[17] In its 2006 10-K filing, in the Introductory section in which eBay formally describes its business strategy, it writes: "We have aggregated a significant number of buyers, sellers, and items listed for sale, which, in turn, has resulted in a vibrant online commerce environment. *Our sellers generally enjoy high conversion rates* and our buyers enjoy an extensive selection of broadly-priced goods and services." (eBay 2006c, emphasis added).

is strictly less than one whenever there is a positive probability of bidders' abstaining from the first auction, i.e., whenever $\underline{v} - \mathbb{E}_{v_2} S(v_2; \overline{v}) < c$ (roughly, the seller's salvage value is high relative to the minimum bidder value). The success rate of the Separate Auctions is always strictly less than one in symmetric PBNE, though there exist asymmetric PBNEs with perfect success rates.

The perfect success rate of course should not be taken literally. The assumption of the model that is unrealistic in practice is the restriction against sellers setting reserve prices in excess of $\underline{v}$.[18]

5 Simulation Results: Magnitude of Effects

On the one hand, eBay made two auction-platform-design decisions that each increased the efficiency of its single-unit auction design (Theorem 1). On the other hand, the eBay market design is not efficient (Remark 1) – only a multi-object auction design can efficiently solve the multi-object allocation problem. So should we admire eBay's clever design or rue its failure to implement simultaneous ascending auctions? The primary purpose of this section is to calculate the magnitudes, for many different distributions of bidder values, of (i) the efficiency gains from Sequencing and Information, and (ii) the remaining inefficiency from not using an efficient multi-object auction. A secondary purpose of this section is to compare revenues across the various auction formats.

5.1 Simulation Methodology

We assume that bidders' values are distributed according to a Beta distribution. The advantages of the Beta distribution are (i) it is parsimonious (it is fully specified by two shape parameters α and β); (ii) it is continuous and differentiable with full support on $[0, 1]$; and (iii) it allows for many different kinds of plausible value distributions, including uniform, $\cup-$shaped, and $\cap-$shaped. A second assumption is that the seller's salvage value, and hence minimum allowable bid, is $c = \underline{v} = 0$.

Given a particular α, β, and n the methodology is as follows. First, I use Zeithammer's (2006) procedure to compute $S(v_2, c_1)$ on a fine grid.[19] I then can use (1) to directly calculate first-round bids for the sequential auction with information hidden,[20] and (2) plus the strict

[18] Meeker (2006) provides a recent estimate of eBay's average success rate as 38%, and states that eBay's success rate had been around 50% in 1998 (its IPO).

[19] Specifically, I use a step size of 0.005 and an error tolerance of 0.001.

[20] Types for whom $v_1 - \mathbb{E}_{v_2} S(v_2, \overline{v}) < c = 0$ will wish to randomize between abstaining and bidding zero. Rather than simulate the randomization, I calculated outcomes under two polar cases: first with all such types abstaining, and second with all such types bidding zero, with the winner determined based on the

monotonicity of $v_1 - S(v_2, b_1) - b_1$ in b_1 to solve for first-round bids for the sequential auction with information revealed.[21] In the second round of the sequential auctions, and in the separate auctions, bidders bid their exact values. Once I have obtained the bidding functions for some (α, β, n) tuple I draw 1,000,000 sets of n bidders' values from $Beta(\alpha, \beta)$ and calculate revenues and surplus for each. The results reported in Table 1 are the averages over these 1,000,000 sets of values.

Additionally, for each set of values I calculate the revenues and surplus from an efficient Vickrey-Clarke-Groves auction (in which bidders have a dominant strategy of stating their true values).

5.2 Simulation Results

The following table summarizes the results of numerical calculations for $n = 3$ and $\alpha, \beta \in \{0.2, 0.5, 0.8, 1, 2, 5\}$.

[Insert Table 1: Efficiency and Revenue Performance]

5.2.1 Efficiency Performance

The main thing to note is that the Sequential Auction with Information Revealed generates 98.9% of the efficient surplus or greater in all specifications.[22] Either separating the auctions or hiding information risks inefficiency an order of magnitude larger.

How can we explain the consistently-strong efficiency performance of the Sequential Auction with Information Revealed?

First, recalling that $\mathbb{E}(V_{1:z})$ denotes the expectation of the 1st-highest of z draws from the bidder value distribution $f(\cdot)$, and recalling that the Sequential Auction with Hidden Information implements the greedy allocation, we can write:

$$(3) \qquad \mathbb{E}(V_{1:n}) + \mathbb{E}(V_{1:(n-1)}) = \text{Surplus(Seq/Hidden)} \underbrace{<}_{\text{Proposition 1}} \text{Surplus (Seq/Revealed)}$$

Second, let K be a random variable distributed Binomial$(n, \frac{1}{2})$, with $\Pr(K = k)$ indicat-

highest first-round value.

The two cases differ only when all n bidders wish to abstain, and so the average difference is quite small. (For the reported simulation parameters, the maximum surplus difference is 0.0009 and the maximum revenue difference is 0.008). Reported results are for the case where all types abstain.

[21] Specifically, bids are calculated as a weighted average of the two grid points between which $v_1 - b_1 - S(v_2, b_1)$ crosses zero.

[22] Note too that the 98.9% is a measurement of the social surplus created from the idiosyncratic portion of bidders' value functions. As an extreme illustration, if bidders' values are distributed Beta on [100,101], then surplus is greater than 99.98% for all specifications.

ing the probability that exactly k of n bidders in the Separate Auctions with Information Revealed enter the first auction (i.e., their first value is higher than their second value). Conditional on k entrants, the highest value in the first auction is the highest of $2k$ random draws from $f(\cdot)$, and so we can write:

$$(4)\quad \mathbb{E}_k[\mathbb{E}(V_{1:2k})+\mathbb{E}(V_{1:2(n-k)})] = \text{Surplus(Sep/Revealed)} \underbrace{<}_{\text{Proposition 2}} \text{Surplus (Seq/Revealed)}$$

Finally, because bidders have unit demand we can provide a simple upper bound to the surplus creation of an efficient VCG auction:

$$(5)\qquad \text{Surplus (Seq/Revealed)} \underbrace{<}_{\text{Remark 1}} \text{Surplus (VCG)} \underbrace{<}_{\text{unit demand}} 2\mathbb{E}(V_{1:n})$$

Combining (3)-(5) we can provide the following bound on the inefficiency of the eBay auction format:

$$\begin{aligned} \text{"eBay Inefficiency"} &= \text{Surplus (VCG) - Surplus (Seq/Revealed)} \\ &< \min\left(\underbrace{\mathbb{E}(V_{1:n}) - \mathbb{E}(V_{1:(n-1)})}_{\text{Inefficiency Bound: Seq/Hdn}}, \underbrace{\mathbb{E}(V_{1:n}) - \mathbb{E}_k(\mathbb{E}(V_{1:2k}))}_{\text{Inefficiency Bound: Sep/Rev}}\right) \end{aligned}$$

The Sequential Auction with Hidden Information squanders a large amount of social surplus when a bidder with a much-higher second-round value than the other bidders also has a marginally-higher first-round value than the other bidders, and so ends up winning the first auction instead. This risk is greatest when $\mathbb{E}(V_{1:n}) - \mathbb{E}(V_{1:(n-1)})$, i.e., the marginal expected value of including the n^{th} bidder in the second auction, is large (large first differences, e.g., because high values are rare).

The Separate Auction with Revealed Information squanders a large amount of social surplus when a bidder with high values for both objects enters the "wrong" auction: the object for which he has a marginally higher value has very strong competition, while the other has very weak competition. This risk is greatest when $\mathbb{E}(V_{1:n}) - \mathbb{E}_k(\mathbb{E}(V_{1:2k}))$ is large, i.e., when the marginal expected value of including the z^{th} bidder in an auction is highly concave in z (large second differences, e.g., because high values are common).

Theorem 1 tells us that the eBay auction format – Sequential Auctions with Information Revealed – performs better than the better of these two alternatives. The insight that explains the eBay format's robust performance is that the performance of these two alternatives is in a sense negatively correlated: distributions of bidder values that yield large values of $\mathbb{E}(V_{1:n}) - \mathbb{E}(V_{1:(n-1)})$ yield small values of $\mathbb{E}(V_{1:n}) - \mathbb{E}_k(\mathbb{E}(V_{1:2k}))$, and vice versa. For any bidder value distribution $f(\cdot)$, the function $\mathbb{E}(V_{1:z})$ of z has to satisfy: (i) decreasing first-differences (the first bidder is more valuable than the second, etc.); and (ii) $\lim_{z\to\infty}\mathbb{E}(V_{1:z}) = \overline{v}$ (values are bounded above by $\overline{v}$). So if $\mathbb{E}(V_{1:z})$ is very concave in z (has large second differences) it must quickly become quite flat in z (has small first-differences).

To see this in Table 1, consider the set of rows with $n = 3$ and $\alpha = .2$. As we increase β the distribution becomes more left skewed, i.e., higher values become rarer and $\mathbb{E}(V_{1:z})$ becomes more flat in z. The result is that the efficiency performance of the Separate Auctions with Information Revealed improves from 89.7% to 98.0%, while the performance of the Sequential Auctions with Information Hidden weakens from 96.0% to 90.1%. By contrast, the performance of the Sequential Auctions with Information Revealed – which is better than the better of these two alternatives – stays pretty constant, at between 98.9% and 99.1%.

5.2.2 Revenue Performance

Referring now to the revenues columns of Table 1, we see that the sequential auction with information revealed generates at least 93.7% of the VCG revenues over the parameters simulated. Revealing information increase total expected revenues over all parameters simulated. Though there exist parameters for which, with information revealed, separating the two auctions increases expected revenues, the overall picture is that the eBay format has the best revenue performance of the taxonomy.

5.2.3 Discussion of Magnitudes

To put the effects of Information and Sequencing in perspective, we note that in a classic paper Riley and Samuelson (1981) calculate that the revenue benefit of utilizing an optimal reserve price in the sale of a single item to two bidders with values uniform on $[0, 1]$ is $\frac{5}{12} - \frac{1}{3} = 0.0833 = 25\%$. Here, the efficiency benefit of Sequencing and Information in the sale of two items to three unit-demand bidders with values drawn from the uniform on $[0, 1]$ is $1.44 - 1.06 = 0.38 = 36\%$, and the revenue benefit is $0.753 - 0.375 = 0.38 = 100\%$. By contrast, the welfare loss from not using an efficient multi-object auction is $1.45 - 1.44 = 0.01 < 1\%$.

5.2.4 Changing the Number of Bidders

Increasing the number of bidders from $n = 3$ improves the efficiency and revenue performance of all auction formats. All formats benefit from the increase in the expectation of the highest and second-highest bidder values in a given auction. The Separate auctions benefit additionally because the risk of an item going unallocated drops quickly in n (the probability of an item going unallocated is $\frac{1}{2^{n-1}}$), and so the difference in performance between the Sequential and Separate auctions gets smaller. For instance, if $n = 8$, all formats but the Sequential Auction with Information Hidden achieve efficiency performance of 98%. See Table 2.

[Insert Table 2: Changing the Number of Bidders]

5.2.5 Changing the Number of Rounds

For robustness, it is desirable to check that the main efficiency results do not depend on the assumption that there are exactly two items auctioned.

Increasing the number of rounds beyond two makes calculating equilibrium behavior for the Sequential Auction with Information Revealed difficult. What makes the case of two rounds tractable is that a bidder's continuation surplus depends on a single auction in which it is a dominant strategy to bid his value. So we can compute $S(v_2, c_1)$ directly: the continuous state variable c_1 conveys information about the set of continuing bidders (see fn. 16), and given this set and v_2 we can calculate expected surplus. If there are $R \geq 3$ rounds, then in all rounds but the last continuation surplus needs to be computed recursively, which is complex because there are continuous value inputs and state variables. An additional problem is that bidders' information sets may be asymmetric after the first round: if, as on eBay, the second-highest bid is disclosed, then the bidder whose bid this was has a different information set than do his opponents.

Thus, to provide this robustness check we will have to make simplifying assumptions. First, suppose that each value is drawn from a discrete value distribution with V possible values. This introduces atoms and gaps, and so departs from the main model, but it reduces the set of possible types to V^R. Second, assume that only the winning bid is disclosed, and not the price this bidder actually paid (i.e., not the second-highest bid). This reduces the number of possible histories (states) to $(V^R)^{R-1}$. Finally, loosely inspired by Pakes and McGuire (1994), restrict attention to equilibria in which a bidders' bid in round r is a symmetric function of (i) the state; (ii) his value in round r; and (iii) the distribution of his values in rounds $r-1...R$; but not the order in which these future values occur. If we set $V = 2$, then this reduces the number of possible equilibrium bids in round r to $2(R-r+1)$, and so the set of possible histories is reduced to $2^{R-1}R!$.

Given these assumptions, for small enough R the set of possible histories and types is small enough that we can solve recursively for equilibrium behavior. (For $R = 2$ we can solve the game analytically; see Appendix B). Results for R up to 7 ($\approx 3x10^5$ histories), and for various values of p, the likelihood that a particular bidder value is the lower of the two possible values, are reported in Table 3.

[Insert Table 3: Robustness Check: Additional Rounds]

Reassuringly, the same qualitative and quantitative features emerge from this analysis. The Sequential Auction with Information Revealed has the strongest efficiency performance of the taxonomy across all parameter values, and always has efficiency of at least 98%.[23] Hiding information is most harmful when high-values are rare ($p = 0.8$), whereas Separating auctions is most harmful when high-values are plentiful ($p = 0.2$). One new pattern that emerges is that the efficiency performance of the separate auctions weakens with market size (holding the buyer-seller ratio constant). This is because the coordination problem amongst bidders with multiple high values is exacerbated.

6 Conclusion

The original hypothesis of my investigation into internet auctions for imperfect substitutes was: (i) eBay auctions are inefficient because they elicit single-dimensional bids from bidders with multi-dimensional value information (Remark 1); (ii) multi-object auctions are efficient; (iii) therefore eBay should use some kind of multi-object auction format; (iv) but market forces have not corrected this failure because of eBay's network externalities; (v) which were won because eBay's auction design *is* efficient for collectibles, its original focus.

This paper does not falsify my original hypothesis per se, but does suggest there's more to the story than Remark 1. The extent to which the lack of simultaneous ascending auctions on the internet should count as a "market failure" depends on the magnitudes of the potential welfare gains. What emerges from this paper is the understanding that the allocative inefficiency of eBay's auction marketplace may be quite small. Theorem 1 indicates that the eBay format is the least inefficient out of a real-world-motivated taxonomy of market designs. The simulation results suggest that the eBay format robustly achieves 99% of the maximum expected surplus, whereas each of the other formats can perform an order of magnitude worse. As Klemperer (2002) and Roth (2002) each emphasize, sometimes it is the tiny details – here, Sequencing and Information – that make the largest difference in market design.

A promising direction for future research is to explore the importance of participation costs in auction-marketplace design. In a sense, this paper handles these costs in a highly

[23] For $R = 2$ the Sequential Auction with Information Revealed is fully efficient when there are two values. There can be inefficient allocations for any $R \geq 3$. For instance, if there are 4 bidders whose types are $\{HHL, HLH, HLH, LLL\}$ and the HHL type wins round one (i.e., he wins the coin toss against the HLH bidders who bid the same as he) then the allocation is inefficient.

reduced form: Separate auctions correspond to a high-cost regime; Sequential auctions to a low-cost regime; and Combinatorial auctions correspond to a regime in which even complex participation is costless. With respect to eBay in particular, there is some evidence that participation costs are quite high (Juda, 2005), and that eBay itself is aware of this: a recent feature, called Bid Assistant, is claimed to save "busy buyers" "valuable time", precisely by making it easier to participate in a sequence of auctions (with information revealed).

A Proofs

Proof of Proposition 1. The argument proceeds in two steps. First we show the result assuming that no bidders abstain from the sequential auction with information hidden (that is, $\underline{v} - \mathbb{E}_{v_2} S(v_2; \overline{v}) > c$). Second, we show that the argument is easily adapted to the case where some types may abstain.

Case 1: No types abstain from the first of two sequential auctions with information hidden.

With future-object information hidden, the equilibrium allocation is greedy: whichever bidder has the highest value for the first object wins the first auction, and whichever remaining bidder has the highest value for the second object wins the second auction. This follows from the independence of values across objects.

To prove that revealing information increases efficiency we only need to look at cases where revealing information alters the allocation, i.e. causes a non-greedy allocation.

Consider the bidder, i, with the highest first-round value: $v_{i1} \geq v_{k1}, \forall k$. Of i's $n-1$ opponents, let bidder j be the one with the highest first-round bid in equilibrium with information revealed: $b_{j1}^{r*} \geq b_{k1}^{r*}, \forall k \neq i$.[24] From (2), we know that for both i and j :

$$b_{i1}^{r*} = v_{i1} - S(v_{i2}, b_{i1}^{r*}) \tag{6}$$

$$b_{j1}^{r*} = v_{j1} - S(v_{j2}, b_{j1}^{r*}) \tag{7}$$

If $b_{i1}^{r*} > b_{j1}^{r*}$, then the allocation is greedy and efficiency has not been affected by the release of future-object information.

If $b_{i1}^{r*} \leq b_{j1}^{r*}$, then revealing information has altered the allocation (in the case of ties, it alters the allocation when j wins the coin toss). We need to show that this alteration is efficiency enhancing in expectation.

[24] In the event of a tie for the role of either i or j (which occur with probability zero), we randomize, taking care to select j after the determination of i.

Define the random variable $v_{K2}(x)$ as the largest second-round value out of $n-2$ bidders randomly drawn from the set of types who bid weakly less than x in the first round. So, given the way we chose j, the largest second-round value of bidders other than i and j is the random variable $v_{K2}(b_{j1}^{r*})$.

We use the following crediting scheme for social surplus creation for the case where revealing information changes the allocation:

	Contribution to Social Surplus	
	Seq / Hidden	Seq / Revealed
i	v_{i1}	$\max(v_{i2} - v_{K2}(b_{j1}^{r*}), 0)$
j	$\max(v_{j2} - v_{K2}(b_{j1}^{r*})), 0)$	v_{j1}
$k \neq i, j$ (combined)	$v_{K2}(b_{j1}^{r*})$	$v_{K2}(b_{j1}^{r*})$

Notice that the crediting scheme treats the two rounds differently. In the first round, the winner (i or j) is credited with his full value. In the second round, the set of bidders $k \neq i, j$, each of whom bids weakly less than b_{j1}^{r*} by construction, is credited with its maximum value. If whichever of i or j loses the first round then has a higher value still, he is credited just with his marginal contribution over and above that of the set $k \neq i, j$. So we need to show that:

(8)
$$\mathbb{E}\left[\max(v_{i2} - v_{K2}(b_{j1}^{r*}), 0) + v_{j1} + v_{K2}(b_{j1}^{r*})\right] \geq \mathbb{E}\left[v_{i1} + \max(v_{j2} - v_{K2}(b_{j1}^{r*}), 0) + v_{K2}(b_{j1}^{r*})\right]$$

By Lemma 2, $S(v_{j2}, b_{j1}^{r*}) = \mathbb{E}[\max(v_{j2} - v_{K2}(b_{j1}^{r*}), 0)]$ and $S(v_{i2}, b_{j1}^{r*}) = \mathbb{E}[\max(v_{i2} - v_{K2}(b_{j1}^{r*}), 0)]$: a bidder's expected profits in a single unit auction against a set of $n-2$ bidders each of whom bids weakly less than b_{j1}^{r*} (the definition of the $S(\cdot)$ function) is equal to his expected marginal contribution to social surplus against that same group. So (8) reduces to:

$$S(v_{i2}, b_{j1}^{r*}) + v_{j1} \geq S(v_{j2}, b_{j1}^{r*}) + v_{i1} \tag{9}$$

Using the equilibrium bids (6) and (7), adding and subtracting $S(v_{i2}, b_{i1}^{r*})$, and rearranging terms, this becomes

$$b_{j1}^{r*} - b_{i1}^{r*} \geq S(v_{i2}, b_{i1}^{r*}) - S(v_{i2}, b_{j1}^{r*}) \tag{10}$$

which obtains since by assumption $b_{j1}^{r*} \geq b_{i1}^{r*}$ and by Lemma 1, $\frac{dS(v_{i2},x)}{dx} > -1$. So for any set of bidders, whenever future-object information induces a non-greedy allocation, surplus is higher in expectation. The efficiency improvement is strict whenever $b_{j1}^{r*} > b_{i1}^{r*}$.

Case 2: Some types abstain from the first of two sequential auctions with information hidden.

If there is at least a single bidder who does not abstain from the first auction, we can use the argument for Case 1 to show that revealing information increases social surplus. The allocation with information hidden is greedy, and we showed above that the allocation with information revealed improves upon the greedy allocation. We need to worry however about the case where all bidders abstain from the first auction, since there will be n rather than $n-1$ participants in the second auction.

Let q be the bidder who submits the highest first-round bid in equilibrium with information revealed: $b_{q1}^{r*} \geq b_{j1}^{r*}, \forall j$ (that such a q exists almost always follows from Zeithammer's (2006; Proposition 1) no abstentions result). Adapting the notation from case 1, we will write $v_{K2}^{n-1}(x)$ for the random variable equal to the largest second-round value out of $n-1$ bidders randomly drawn from the set of types who bid weakly less than x.

We credit the $n-1$ bidders other than q, each of whom would bid weakly less than b_{q1}^{r*} with information revealed, with $v_{K2}^{n-1}(b_{q1}^{r*})$ of surplus for either information environment. With information hidden, then by construction q also abstains, and we credit him with $\max(v_{q2} - v_{K2}^{n-1}(b_{q1}^{r*}), 0)$ of surplus. With information revealed, q wins the first auction, and we credit him with v_{q1} of surplus creation. Using (2) and Lemma 2 we have:

$$v_{q1} = b_{q1}^{r*} + S(v_{q2}, b_{q1}^{r*}) = b_{q1}^{r*} + \mathbb{E}(\max(v_{q2} - v_{K2}(b_{q1}^{r*}), 0)) \tag{11}$$

Since the largest of $n-1$ draws from a distribution is larger than the largest of $n-2$ draws from the same distribution $\mathbb{E}(\max(v_{q2} - v_{K2}(b_{q1}^{r*}), 0) \geq \mathbb{E}(\max(v_{q2} - v_{K2}^{n-1}(b_{q1}^{r*}), 0))$. Zeithammer's (2006) Proposition 1 gives that b_{q1}^{r*} is weakly greater than the seller's salvage value. So for any set of n bidders who abstain if information is hidden, revealing information increases expected surplus, as required. □

Proof of Proposition 2. First, note that if information is hidden, then any equilibrium of the separate auctions must be less efficient than the equilibrium of the sequential auctions. The two sequential auctions have n and $n-1$ participants, whereas the two separate auctions have just $n-k$ and k for some $k = 1, ..., n-1$, and in any of these auctions the participating bidder with the highest value wins.

What remains is to show that the sequential auction with information revealed generates more surplus than the separate auctions with information revealed. Let i be the bidder who wins object one when the auction format is Seq / Revealed. We consider three cases.

Case 1: i also wins object one under Sep / Revealed.

Clearly the Seq / Revealed generates weakly more welfare than the Sep / Revealed in this case because the set of participants in the auction for the second object under Sep / Revealed is a subset of the set of participants for Seq / Revealed, and in either auction format the bidder with the highest valuation wins the object in equilibrium.

Case 2: i enters the auction for object one under Sep / Revealed, but some other bidder j wins object one.

This argument is simple once we define an appropriate scheme for crediting social surplus creation to bidders (similar to that used in the proof of Proposition 1).

We know that each of the bidders $k \neq i, j$ bids weakly less than b_{i1}^{r*} in the Seq / Revealed. Credit this set of bidders with surplus $v_{K2}(b_{i1}^{r*})$ in the Seq / Revealed , and with $v_{K2}^{Sep}(b_{i1}^{r*})$ in the Sep / Revealed, where $v_{K2}^{Sep}(x)$ is the random variable equal to the largest second-item value out of $n-2$ bidders randomly drawn from the set of types who (1) bid weakly less than x in the first round, and also (2) value the second object more highly than the first.

Now credit social surplus according to the following table:

	Contribution to Social Surplus	
	Seq / Revealed	Sep / Revealed
i	v_{i1}	0
j	$\max(v_{j2} - v_{K2}(b_{i1}^{r*}), 0)$	v_{j1}
$k \neq i, j$ (combined)	$v_{K2}(b_{i1}^{r*})$	$v_{K2}^{Sep}(b_{i1}^{r*})$

Using Lemma 2 and the observation that clearly $v_{K2}(b_{i1}^{r*}) \geq v_{K2}^{Sep}(b_{i1}^{r*})$, what we need to show is:

$$v_{i1} + S(v_{j2}, b_{i1}^{r*}) \geq v_{j1} \tag{12}$$

Adding and subtracting both $S(v_{i2}, b_{i1}^{r*})$ and $S(v_{j2}, b_{j1}^{r*})$ and using the equilibrium bids for i and j gives

$$b_{i1}^{r*} + S(v_{i2}, b_{i1}^{r*}) - b_{j1}^{r*} \geq S(v_{j2}, b_{j1}^{r*}) - S(v_{j2}, b_{i1}^{r*}) \tag{13}$$

which inequality is seen to obtain using Lemma 1, $b_{i1}^{r*} \geq b_{j1}^{r*}$ and the observation that $S(v_{i2}, b_{i1}^{r*}) \geq 0$. The efficiency gain is strict whenever $b_{i1}^{r*} > b_{j1}^{r*}$.

Case 3: i enters the auction for object two under Sep / Revealed, and so some other bidder j wins object one.

We use a counting scheme similar to that for the previous case; indeed all of the cells of the surplus crediting scheme are the same but for one. The difference is i's contribution to social surplus in Sep / Revealed, which now will be positive since he enters the second auction. This cell is now equal to $\max(v_{i2} - v_{K2}^{Sep}(b_{i1}^{r*}), 0)$.

Since $v_{K2}^{Sep}(b_{i1}^{r*}) \leq v_{K2}(b_{i1}^{r*})$ for any realized set of bidders $k \neq i, j$, it follows that $\mathbb{E}[\max(v_{i2} - v_{K2}^{Sep}(b_{i1}^{r*}), 0) + v_{K2}^{Sep}(b_{i1}^{r*})] \leq \mathbb{E}[\max(v_{i2} - v_{K2}(b_{i1}^{r*}), 0) + v_{K2}(b_{i1}^{r*})]$. Using Lemma 2, this reduces the problem to showing that $S(v_{i2}, b_{i1}^{r*}) + v_{j1} \leq S(v_{j2}, b_{i1}^{r*}) + v_{i1}$. Add and subtract $S(v_{j2}, b_{j1}^{r*})$ and use the equilibrium bid formulas, and we have $b_{i1}^{r*} - b_{j1}^{r*} > S(v_{j2}, b_{i1}^{r*}) - S(v_{j2}, b_{j1}^{r*})$ which is true by Lemma 1. The efficiency gain is strict whenever $b_{i1}^{r*} > b_{j1}^{r*}$.

Since these 3 cases are exhaustive, the proof is complete. □

Proof of Lemma 3. Lemma 2 and the symmetry of bidders implies $0 < \frac{\partial S(v_2,c_1)}{\partial v_2} \leq 1$ for $v_2 > \underline{v}$ because such bidders win the second auction with strictly positive probability. Total differentiation of (2) with respect to v_2, using Lemma 1, then yields $\frac{db_1^{r*}(v_1,v_2)}{dv_2} < 0$ for $v_2 > \underline{v}$, with weak inequality for $v_2 = \underline{v}$.

Pick an arbitrary v_1. Define $\underline{v_2}(v_1, c_1) \geq \underline{v}$ as the least v_2 for which $b_1^{r*}(v_1, v_2) \leq c_1$. Now define

$$F(v_2|c_1, v_1) = \frac{\int_{\min(\underline{v_2}(v_1,c_1),v_2)}^{v_2} f(x)dx}{\int_{\underline{v_2}}^{1} f(x)dx}$$

It is easy to see that $F(v_2|c_1, v_1) \leq F(v_2|\overline{v}, v_1)$, and so $\int_{\underline{v}}^{\overline{v}} F(v_2|c_1, x)f(x)dx \equiv F(v_2|c_1) \leq F(v_2|\overline{v})$, i.e., the distribution of a second-round opponent's bid given a first-round winning bid of c_1 first-order stochastically dominates that given a first-round winning bid of $\overline{v}$. From this it follows that the distribution of the *maximum* second-round opponent bid given c_1 first-order stochastically dominates that given $\overline{v}$, and so a bidder with value v_2 has a greater expected contribution to social surplus in the second round given c_1 than given $\overline{v}$. Hence, by Lemma 2, $S(v_2, c_1) \leq S(v_2, \overline{v})$. □

Proof of Proposition 4. First, we consider bidders for whom $v_1 - \mathbb{E}_{v_2} S(v_{2,}\overline{v}) \geq c$; such bidders never abstain. With information hidden, these bidders all shade by a common amount $\mathbb{E}_{v_2} S(v_2; \overline{v})$. We need to show that, for any v_1:

$$b_1^{h*}(v_1, -) = v_1 - \mathbb{E}_{v_2} S(v_{2,}\overline{v}) \leq v_1 - \mathbb{E}_{v_2} S(v_{2,} b_1^{r*}(v_1, v_2)) = \mathbb{E}_{v_2} b_1^{r*}(v_1, v_2) \tag{14}$$

That is, the first round bid of a bidder with arbitrary first-round value v_1 with information hidden is less than the expectation of this same bidder's first-round bid with information revealed over all possible realizations of v_2. Since $b_1^{r*}(v_1, v_2) \leq \overline{v}$, the result follows directly from Lemma 3.

Bidders for whom $v_1 - \mathbb{E}_{v_2} S(v_{2,}\overline{v}) < c$ either always abstain or randomize between abstaining and bidding exactly c when information is hidden. Zeithammer (2006, Proposition 1) shows that with information revealed almost all of these bidders bid strictly greater than c. The zero-measure set for whom $v_1 = \underline{v}$ abstains in either information environment. □

B A Simple Two-Value Example

To build intuition about the Sequential Auction with Information Revealed we provide a simple two-value example that is solvable in closed form.

All is as described in Section 1, except let $f(\cdot)$ be a probability mass function with $f(L) = p$ and $f(H) = 1 - p$, and $\underline{v} < L < H < \overline{v}$. Note that the two-value model is not a special case of the main model, because it has atoms and gaps.

There are four types of bidders, $(L, L), (L, H), (H, L)$, and (H, H). First note that (L, L) and (H, L) types will shade their first-round bids by zero, since their second-round profits are zero whatever the outcome of the first round.

$(H, H)'$s bid will only be relevant at the margin against other (H, H) bidders, in the case where there also are not any (H, L) bidders. If there is at least one (L, H) bidder or at least one additional (H, H) bidder, his second round profits if he loses the first round are zero. Thus we can explicitly calculate (H, H)'s unique symmetric PBNE first-round bid as:[25]

$$(15) \qquad b_1^{r*}(H, H) = H - \frac{(n-1)[(1-p)^2][p^2]^{n-2}}{[1-p(1-p)]^{n-1} - p^{n-1}}(H - L) = H - S(H, b_1^{r*}(H, H))$$

Finally, consider type (L, H). His strategy is surprising. We might expect type (L, H) to bid quite cautiously in the first round, since his future is attractive and his present bleak. But if he bids any amount less than L, his bid is only relevant at the margin when he faces exclusively (L, H) opponents, i.e., when his future is actually worth zero. Relative to unconditionally losing, *marginally* losing as an (L, H) conveys especially bad news about second-round profit opportunities: $S(H, x) = 0$ for any $x < L$. Consequently, for equilibrium to exist in this model we need to assume that there exists an L^- strictly less than L such that bid amounts in the interval (L^-, L) are not permitted. The (L, H) types bid L^- in equilibrium.[26]

Example 1 *Let* $L = 10$, $H = 20$, $n = 3$, *and* $p = 1/2$. *Then* $\mathbb{E}_{v_2} S(v_2, H) = 2.5$ *and* $S(H, b_1^{r*}(H, H)) = 4$. *That is, an* (H, H) *bidder is indifferent between marginally winning and marginally losing the first auction at his equilibrium bid of* 16.

Efficiency and Revenues It follows directly from the bidding functions that for the two-value model the eBay auction format is efficient, and it is easy to see that each of the other

[25] This can be found by noting that $b_1^{r*}(H, H) = H - \Pr(Y|Z)(H - L) = H - \frac{\Pr(Y)}{\Pr(Z)}(H - L)$, where:
Y ="I face exactly one (H, H) opponent and $n - 2$ (L, L) opponents: I want to lose Round 1"
Z ="I face at least one (H, H) opponent and no (H, L) opponents: My Round 1 bid matters"

[26] The analogous result without atoms in the bidder value distribution is that a zero-measure set of bidders – those whose first-round values are equal to $\underline{v}$ – may abstain from the first auction (Zeithammer 2006b; Proposition 1). Also note that if $L - L^-$ is large enough then the (L, H) types may randomize.

auction formats can be inefficient:

Proposition A1 *For the two-value model, the unique symmetric PBNE of the Sequential Auction with Information Revealed is efficient. Any PBNE of any of the alternate formats in the taxonomy is inefficient.*

The first-round competition effect is surprisingly large in the two-value model. (L, L) and (H, L) types bid their full value instead of shading, and (L, H) types shade by the minimum possible amount. There even exist parameter values for which (H, H) types with Information Revealed bid more than $(H, -)$ types with Information Hidden! For example, if $n = 3$, $p = 0.25$, then $S(H, b_1^{r*}(H, H)) = 11.8\% * (H - L)$, whereas $\mathbb{E}_{v_2} S(v_2, \overline{v}) = 18.75\% * (H - L)$. (This effect obtains for p low enough or n high enough.)[27]

Proposition A2 *(Magnitude of First-Round Competition Effect) For the two-value model, for low enough p or high enough n, revealing information in a sequence of two auctions unambiguously increases all types' first-round bids, in all states of the world.*

References

[1] Athey, Susan and Glenn Ellison, 2008. "Position Auctions with Consumer Search." Working Paper.

[2] Athey, Susan, Jonathan Levin and Enrique Seira, 2004. "Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions." Working Paper.

[3] Bajari, Patrick and Ali Hortacsu, 2003. "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions." *RAND Journal of Economics*, 34(2), 329-55.

[4] Bajari, Patrick and Ali Hortacsu, 2004. "Economic Insights from Internet Auctions: A Survey." *Journal of Economic Literature*, 42(2), 457-86.

[27] The fact that *all* types always bid more relies on the atom at (H, H). In the two-value model, when an (H, H) bidder is marginal, he does not know if he is tied at the margin with one additional (H, H) bidder, in which case his second round profits might be $H - L$, or with multiple (H, H) bidders, in which case his second round profits will be zero. Suppose that instead of the atom at (H, H) there is $(1 - p)^2$ of density in the region $[H - \delta, H + \delta]^2$. Consider type $(H - \delta, H + \delta)$: his first-round bid is lowest out of all "(H, H)" types, and in particular, if his first-round bid is marginal, then the probability that there are other "(H, H)" types who will continue to the second-round is zero. So he will shade his first-round bid by approximately $S(H, \overline{v}) = p^{n-2}(H - L)$, which is more than he would have shaded by if information were hidden. In the limit when $\delta = 0$, he instead bids according to (15) – the difference is that now when he is marginal he may face other (H, H) types in the second round.

[5] Beggs, Alan and Kathryn Graddy, 1997. "Declining Values and the Afternoon Effect: Evidence from Art Auctions." *RAND Journal of Economics*, 28(3), 544-65.

[6] Budish, Eric and Lisa Takeyama, 2001. "Buy Prices in Online Auctions. Irrationality on the Internet?" *Economics Letters*, 72, 325-33.

[7] Bulow, Jeremy and Paul Klemperer, 1996. "Auctions versus Negotiations." American Economic Review, 86(1), 180-94.

[8] Bulow, Jeremy and Paul Klemperer, 2007. "When are Auctions Best?" Working Paper.

[9] Cohen, Adam, 2002. The Perfect Store: Inside eBay. Little, Brown and Company.

[10] Cohen, Jane and Bob Grossweiner, 2006. "Industry Profile: Jeff Fluhr." Celebrity Access, October 9, accessed on www.stubhub.com.

[11] Cramton, Peter, Yoav Shoham and Richard Steinberg (eds.), 2006. Combinatorial Auctions. The MIT Press.

[12] eBay, 1998. eBay Inc. form S-1, July 15.

[13] eBay, 2006a. "2006 Analyst Day Presentation." Accessed at http://investor.ebay.com/downloads/AnalystDay_2006.pdf.

[14] eBay, 2006b. "eBay Inc. Announces Fourth Quarter and Full-Year 2005 Financial Results." Accessed at http://investor.ebay.com/financial.cfm.

[15] eBay, 2006c. eBay Inc. form 10-K. Accessed at http://investor.ebay.com/sec.cfm.

[16] Ellison, Glenn, Drew Fudenberg and Markus Mobius, 2004. "Competing Auctions." *Journal of the European Economic Association*, 2(1), 30-66.

[17] Engelbrecht-Wiggans, Richard, 1993. "Optimal Auctions Revisited." *Games and Economic Behavior*, 5, 227-39.

[18] Engelbrecht-Wiggans, Richard, 1994. "Sequential Auctions of Stochastically Equivalent Objects." *Economics Letters*, 44, 87-90.

[19] Eyster, Erik and Matthew Rabin, 2005. "Cursed Equilibrium" *Econometrica*, 73(5), 1623-72.

[20] Hagiu, Andrea and Bruno Jullien, 2007. "Designing a Two-Sided Platform: When to Increase Search Costs?" Working Paper.

[21] Harrington, Richard, 2006. "Ticket Auction Trend May Cost You." Washington Post, June 2, WE27.

[22] Hidvegi, Zoltan, Wenli Wang and Andrew B. Whinston, 2006. "Buy-price English Auction." *Journal of Economic Theory*, 129(1), 31-56.

[23] Juda, Adam I., 2005. "The Sequential Auction Problem: An Analysis and Solution." Harvard University Working Paper.

[24] Kaiser, Laura Fisher and Michael Kaiser, 1999. The Official eBay Guide to Buying, Selling, and Collecting Just About Anything. Simon & Schuster.

[25] Klemperer, Paul, 2002. "What Really Matters in Auction Design." *Journal of Economic Perspectives*, 16, 169-89.

[26] Lee, Hanh Ah and Ulrike Malmendier, 2006. "Do Consumers Know their Willingness to Pay? Evidence from eBay Auctions." Stanford University Working Paper.

[27] Lewis, Gregory, 2006. "Asymmetric Information, Adverse Selection, and Seller Reputation on eBay Motors." Working Paper.

[28] Lewis, Gregory, 2008. "What's in a Name? Online versus Offline Repuations." In Preparation.

[29] McAfee, Preston and John McMillan, 1987. "Auctions with Entry." *Economics Letters*, 23, 343-7.

[30] Meeker, Mary, 2006. "eBay: Listings Data – Reiterate Overweight." Morgan Stanley Research, May 18.

[31] Milgrom, Paul, 2004. Putting Auction Theory to Work. Cambridge University Press.

[32] Milgrom, Paul and Robert Weber, 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica*, 50(5), 1089-1122

[33] Miniter, Paulette, 2008. "Bidz.com Battles to Restore Credibility." Accessed at http://www.smartmoney.com/undertheradar/ index.cfm?story=20080320-Bidz&afl=yahoo&hpadref=1

[34] Pakes, Ariel and Paul McGuire, 1994. "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model." *RAND Journal of Economics*, 25(4), 555-589.

[35] Peters, Michael and Sergei Severinov, 2006. "Internet Auctions with Many Traders." Journal of Economic Theory, 130(1), 220-245.

[36] Riley, John G. and William F. Samuelson, 1981. "Optimal Auctions." *American Economic Review*, 71, 381-92.

[37] Rochet, Jean-Charles and Jean Tirole, 2006. "Two-Sided Markets: A Progress Report." *Rand Journal of Economics*, 37(3), 645-67.

[38] Roth, Alvin E., 2002. "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics." *Econometrica*, 70(4), 1341-78.

[39] Roth, Alvin E. and Axel 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.

[40] Smith, Ethan, 2006. "Ticketmaster Adapts to Rivals' Online Threat." Wall Street Journal, Sept 12.

[41] Vickrey, William, 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, 16(1), 8-37.

[42] Vickrey, William, 1962. "Auction and Bidding Games." In Recent Advances in Game Theory. The Princeton University Conference, 15-27.

[43] Weber, Robert J., 1983. "Multi-Object Auctions." In Engelbrecht-Wiggans, Richard et al (eds.) Auctions, Bidding and Contracting. New York University Press, 165-91.

[44] Zeithammer, Robert, 2006. "Sequential Auctions with Information about Future Goods." Working Paper.

Table 1. Efficiency and Revenue Performance of Market Designs for Allocating Imperfect Substitutes

				Social Surplus					Revenues				
Simulation Parameters				Multi-Object	Sequential		Separate		Multi-Object	Sequential		Separate	
n	alpha	beta		VCG or SAA	Revealed	Hidden	Revealed	Hidden	VCG or SAA	Revealed	Hidden	Revealed	Hidden
3	0.2	0.2		1.648	99.1%	96.0%	89.7%	69.1%	0.724	94.1%	82.0%	85.2%	44.8%
3	0.2	0.5		1.114	98.9%	94.3%	94.5%	64.0%	0.308	94.0%	74.7%	100.3%	43.3%
3	0.2	0.8		0.829	98.9%	93.0%	96.2%	62.0%	0.191	94.3%	69.8%	102.2%	41.9%
3	0.2	1		0.708	98.9%	92.5%	96.7%	61.2%	0.152	94.5%	67.9%	102.6%	41.1%
3	0.2	2		0.405	99.0%	91.1%	97.6%	59.5%	0.074	95.5%	63.9%	102.8%	39.8%
3	0.2	5		0.177	99.1%	90.1%	98.0%	58.4%	0.029	98.1%	61.4%	102.9%	39.2%
3	0.5	0.2		1.891	99.5%	97.9%	87.5%	76.5%	1.326	94.7%	85.4%	63.9%	44.8%
3	0.5	0.5		1.553	99.1%	96.9%	90.4%	71.0%	0.753	93.9%	81.9%	78.9%	46.2%
3	0.5	0.8		1.282	99.0%	96.3%	92.2%	68.8%	0.521	93.7%	78.6%	84.8%	46.4%
3	0.5	1		1.143	99.1%	96.0%	93.0%	67.8%	0.433	93.7%	77.0%	86.7%	46.2%
3	0.5	2		0.734	99.1%	94.8%	94.6%	65.6%	0.234	93.9%	72.1%	90.3%	45.5%
3	0.5	5		0.350	99.2%	93.7%	95.5%	63.9%	0.098	94.1%	68.4%	91.8%	44.8%
3	0.8	0.2		1.941	99.6%	98.6%	87.3%	79.7%	1.568	95.7%	88.4%	58.0%	45.4%
3	0.8	0.5		1.710	99.3%	97.8%	89.0%	74.7%	1.047	94.6%	85.7%	68.7%	46.4%
3	0.8	0.8		1.486	99.2%	97.3%	90.5%	72.4%	0.781	94.3%	83.3%	74.1%	46.8%
3	0.8	1		1.360	99.2%	97.1%	91.2%	71.5%	0.668	94.3%	82.0%	76.3%	47.0%
3	0.8	2		0.940	99.2%	96.3%	92.9%	69.0%	0.389	94.2%	77.5%	80.9%	46.9%
3	0.8	5		0.481	99.3%	95.4%	94.1%	67.1%	0.173	94.3%	73.5%	83.7%	46.5%
3	1	0.2		1.956	99.7%	98.9%	87.3%	81.0%	1.655	96.1%	89.9%	56.2%	45.8%
3	1	0.5		1.767	99.4%	98.1%	88.6%	76.4%	1.184	95.0%	87.3%	64.8%	46.5%
3	1	0.8		1.569	99.3%	97.7%	89.9%	74.1%	0.915	94.8%	85.5%	69.8%	47.0%
3	1	1		1.453	99.3%	97.5%	90.5%	73.2%	0.795	94.7%	84.4%	71.8%	47.2%
3	1	2		1.042	99.3%	96.8%	92.1%	70.6%	0.484	94.5%	80.4%	76.7%	47.3%
3	1	5		0.554	99.3%	96.0%	93.3%	68.5%	0.223	94.5%	76.2%	79.9%	47.0%
3	2	0.2		1.981	99.9%	99.5%	87.4%	84.0%	1.832	97.6%	93.9%	52.8%	47.4%
3	2	0.5		1.884	99.6%	98.9%	87.9%	80.8%	1.534	96.6%	92.0%	57.2%	47.3%
3	2	0.8		1.763	99.5%	98.6%	88.6%	79.0%	1.315	96.3%	91.0%	60.1%	47.4%
3	2	1		1.683	99.5%	98.5%	89.0%	78.1%	1.201	96.1%	90.4%	61.5%	47.6%
3	2	2		1.349	99.4%	98.0%	90.2%	75.6%	0.843	95.9%	88.2%	65.5%	47.9%
3	2	5		0.824	99.4%	97.3%	91.4%	73.2%	0.449	95.7%	84.8%	69.1%	47.9%
3	5	0.2		1.993	99.9%	99.8%	87.4%	86.0%	1.935	98.9%	97.2%	51.0%	48.8%
3	5	0.5		1.954	99.8%	99.5%	87.6%	84.4%	1.800	98.3%	96.1%	52.7%	48.5%
3	5	0.8		1.900	99.8%	99.4%	87.9%	83.4%	1.681	98.0%	95.5%	54.0%	48.5%
3	5	1		1.860	99.7%	99.3%	88.1%	82.8%	1.611	97.9%	95.2%	54.6%	48.5%
3	5	2		1.665	99.7%	99.0%	88.7%	81.1%	1.338	97.6%	94.1%	56.8%	48.5%
3	5	5		1.233	99.6%	98.5%	89.5%	78.8%	0.896	97.3%	92.4%	59.5%	48.6%
			Max	1.993	99.9%	99.8%	98.0%	86.0%	1.935	98.9%	97.2%	102.9%	48.8%
			Median	1.469	99.3%	97.4%	90.3%	73.2%	0.788	94.7%	84.6%	70.8%	46.9%
			Min	0.177	98.9%	90.1%	87.3%	58.4%	0.029	93.7%	61.4%	51.0%	39.2%

Table 2. Efficiency and Revenue Performance - Effect of Changing the Number of Bidders

Simulation Parameters			Social Surplus					Revenues				
			Combinatorial	Sequential		Separate		Combinatorial	Sequential		Separate	
n	alpha	beta	VCG	Revealed	Hidden	Revealed	Hidden	VCG	Revealed	Hidden	Revealed	Hidden
3	1	1	1.453	99.3%	97.5%	90.5%	73.2%	0.795	94.7%	84.4%	71.8%	47.2%
4	1	1	1.578	99.4%	98.6%	93.9%	77.8%	1.094	97.2%	92.9%	76.4%	52.4%
5	1	1	1.646	99.7%	99.2%	95.9%	81.4%	1.268	98.1%	95.9%	82.8%	59.3%
6	1	1	1.702	99.7%	99.4%	97.1%	84.4%	1.385	98.7%	97.2%	87.9%	64.6%
7	1	1	1.737	99.8%	99.5%	98.0%	86.2%	1.467	99.0%	98.0%	91.2%	69.5%
8	1	1	1.770	99.8%	99.6%	98.5%	87.8%	1.533	99.3%	98.6%	93.6%	73.1%
3	0.5	0.5	1.553	99.1%	96.9%	90.4%	71.0%	0.753	93.9%	81.9%	78.9%	46.2%
4	0.5	0.5	1.695	99.4%	98.4%	93.3%	76.1%	1.131	96.5%	91.6%	78.4%	48.9%
5	0.5	0.5	1.777	99.5%	98.9%	95.2%	80.5%	1.367	97.5%	94.9%	82.5%	54.3%
6	0.5	0.5	1.831	99.7%	99.3%	96.5%	83.4%	1.508	98.4%	96.8%	86.4%	60.5%
7	0.5	0.5	1.867	99.7%	99.5%	97.5%	86.1%	1.614	98.8%	97.6%	89.9%	65.2%
8	0.5	0.5	1.894	99.8%	99.7%	98.2%	88.0%	1.689	99.1%	98.3%	92.1%	69.1%
3	2	2	1.349	99.4%	98.0%	90.2%	75.6%	0.843	95.9%	88.2%	65.5%	47.9%
4	2	2	1.446	99.6%	98.8%	94.2%	80.2%	1.069	97.8%	94.4%	74.6%	56.0%
5	2	2	1.506	99.7%	99.2%	96.2%	83.3%	1.191	98.7%	96.7%	83.2%	64.0%
6	2	2	1.554	99.7%	99.4%	97.5%	85.7%	1.279	99.0%	97.8%	88.5%	69.7%
7	2	2	1.588	99.8%	99.5%	98.3%	87.4%	1.342	99.2%	98.3%	91.9%	74.0%
8	2	2	1.615	99.8%	99.6%	98.6%	88.9%	1.389	99.4%	98.7%	94.2%	77.6%
3	0.5	2	0.734	99.1%	94.8%	94.6%	65.6%	0.234	93.9%	72.1%	90.3%	45.5%
4	0.5	2	0.852	99.3%	97.0%	95.7%	68.1%	0.375	96.7%	84.9%	86.9%	45.8%
5	0.5	2	0.943	99.5%	98.1%	96.5%	71.2%	0.494	97.9%	91.1%	88.8%	49.6%
6	0.5	2	1.010	99.6%	98.5%	97.4%	72.9%	0.587	98.3%	93.9%	90.4%	52.6%
7	0.5	2	1.071	99.6%	98.9%	98.0%	74.7%	0.664	98.7%	95.5%	92.3%	56.0%
8	0.5	2	1.119	99.7%	99.1%	98.2%	76.0%	0.730	98.8%	96.5%	93.5%	58.9%
3	2	0.5	1.884	99.6%	98.9%	87.9%	80.8%	1.534	96.6%	92.0%	57.2%	47.3%
4	2	0.5	1.929	99.8%	99.5%	93.1%	87.0%	1.748	98.4%	96.8%	70.3%	59.9%
5	2	0.5	1.952	99.9%	99.7%	96.1%	91.0%	1.837	99.0%	98.3%	80.4%	70.4%
6	2	0.5	1.964	99.9%	99.8%	97.7%	93.5%	1.882	99.4%	99.0%	87.7%	78.0%
7	2	0.5	1.973	99.9%	99.9%	98.7%	95.3%	1.913	99.6%	99.3%	92.2%	83.2%
8	2	0.5	1.979	99.9%	99.9%	99.2%	96.4%	1.933	99.6%	99.5%	95.2%	87.3%

Table 3. Robustness Check for Efficiency Performance: More Rounds

Simulation Parameters			Social Surplus				
			Multi-Object	Sequential		Separate	
n	R	p	VCG or SAA	Revealed	Hidden	Revealed	Hidden
3	2	0.2	19.805	100.0%	98.6%	86.6%	79.1%
3	2	0.5	17.033	100.0%	95.4%	88.6%	67.9%
3	2	0.8	9.260	100.0%	91.5%	96.8%	58.5%
4	3	0.2	29.946	99.3%	98.5%	80.1%	71.1%
4	3	0.5	27.471	99.6%	93.7%	81.6%	56.5%
4	3	0.8	16.357	99.9%	90.3%	93.3%	44.2%
5	3	0.2	29.990	99.9%	99.7%	86.5%	78.7%
5	3	0.5	28.867	99.8%	96.4%	85.2%	62.1%
5	3	0.8	19.094	99.9%	92.9%	92.5%	45.9%
6	4	0.2	39.996	100.0%	99.8%	82.1%	73.8%
6	4	0.5	39.244	99.7%	96.1%	81.4%	56.1%
6	4	0.8	27.706	99.6%	91.9%	88.9%	38.4%
7	5	0.2	50.000	100.0%	99.8%	79.0%	70.5%
7	5	0.5	49.544	99.3%	96.2%	78.6%	52.7%
7	5	0.8	37.013	99.2%	91.2%	86.0%	33.6%
8	5	0.2	50.000	100.0%	100.0%	83.3%	75.2%
8	5	0.5	49.786	99.8%	98.0%	82.5%	57.2%
8	5	0.8	39.828	99.2%	92.5%	86.1%	35.0%
9	6	0.2	60.000	100.0%	100.0%	80.6%	72.4%
9	6	0.5	59.874	99.7%	98.2%	80.2%	54.3%
9	6	0.8	49.824	98.8%	91.8%	83.5%	31.7%
10	7	0.2	70.000	100.0%	100.0%	78.6%	70.3%
10	7	0.5	69.928	99.7%	98.3%	78.4%	52.4%
10	7	0.8	60.177	98.5%	91.3%	81.2%	29.3%
11	7	0.2	70.000	100.0%	100.0%	81.7%	73.7%
11	7	0.5	69.965	99.8%	99.2%	81.5%	55.8%
11	7	0.8	62.481	98.3%	92.5%	82.1%	30.7%

Notes: Calculations for two-value distribution with High = 10, Low = 0, and seller's cost c << 0
n = number of bidders; R = number of rounds; p = probability that each value is Low