An Efficiency Ranking of Markets Aggregated from Single-Object Auctions

Eric Budish   Robert Zeithammer

July 26, 2016

Abstract: In a wide variety of economic contexts, collections of single-object auctions are used to allocate multiple imperfectly substitutable objects. This paper studies how such collections of auctions aggregate into a multi-object auction market. We identify two features of auction-market design that enhance expected market efficiency: the individual auctions should be conducted in sequence, and information about all of the objects in the sequence should be revealed up front to the bidders. We then show that an auction market with auctions sequenced and information revealed approximates full efficiency in that it is more efficient than the fully efficient VCG mechanism with one fewer bidder.

Keywords: auctions, market sesign, sequential auctions, information in auctions, e-commerce, procurement, platforms, user-interface design, Bulow-Klemperer theorem

---

Contact: Eric Budish (eric.budish@chicagobooth.edu) , Robert Zeithammer (rzeitham@ucla.edu). Acknowledgments: This paper combines and supersedes Zeithammer (2005) and Budish (2008). We are grateful to Susan Athey, Sushil Bikhchandani, Lucas Coffman, Itay Fainmesser, Drew Fudenberg, Ali Hortacsu, Sergei Izmalkov, Ian Jewitt, Paul Klemperer, Robin Lee, Greg Lewis, John Lewis, Ariel Pakes, Robert Ritz, Al Roth, Laura Serban, Birger Wernerfelt, and seminar participants at UT Austin, CERGE-EI, CU Boulder, Harvard, and UCLA for comments and suggestions. We thank Ryan Dorow for excellent research assistance.

1 Introduction

In a wide variety of economic contexts, collections of single-object auctions are used to allocate multiple imperfectly substitutable objects. One important example is government procurement (e.g., highway construction): typically, each contract is allocated via its own sealed-bid auction, even though contracts are likely to be substitutable for suppliers due to capacity constraints. Another well-known example is the online auctioneer eBay, which uses the same single-object auction design for goods with substitutes, such as used automobiles and event tickets, as it does for the one-of-a-kind collectibles for which it was originally designed. Other examples include real-estate auctions, wine auctions, timber auctions, and offline wholesale used-auto auctions.

This paper studies how such collections of single-object auctions "aggregate up" into what we will call an *auction market*. We focus on two prominent market-design features: the timing of the individual auctions, and the amount of information about the objects sold available to bidders in the beginning of the game. The timing of the auctions can be either simultaneous, forcing each bidder to choose which auction to participate in, or sequential whereby all bidders can easily participate in every auction. There are many examples of sequential auction markets in the real world, from government procurement to the way traditional auction houses like Sotheby's and Christie's sell a particular lot of art. Simultaneous markets occur when multiple separate auction markets compete with each other, or when the underlying sequence of auctions is disguised, such as in eBay's new default sort of a consumer's search results. Interestingly, a market for one type of object does not necessarily use the same timing everywhere. For example, wholesale used-car auctions are also held sequentially at the Suwon market in Korea while they are held at least partially simultaneously at the Manheim market in the United States. Real-world auction markets also differ in terms of the information available to bidders at the beginning of the market: both government-procurement and art auctioneers provide detailed information about all of their upcoming auctions. In contrast, industrial procurement auctions are often held with only an expectation of an uncertain future opportunity for the suppliers to win another contract from a different client. We show that both sequencing the auctions and providing information about all objects sold enhances the expected allocative efficiency of the market as a whole. We then show that the market with auctions sequenced and information about future objects revealed, although not fully efficient, in a certain sense approximates full efficiency.

Our model is simple: we consider two objects for sale to $N \geq 3$ potential buyers, with each sale conducted by a second-price sealed-bid auction. Bidders regard the two objects as *imperfect* substitutes in the sense that they have different private valuations of each object and unit capacity; for example, a typical consumer buying a used automobile on eBay Motors values different cars differently and can garage at most one car. Our main technical contribution is a an equilibrium characterization of the sequential auction with information about future objects revealed, generalizing the analyses of Gale and Hausch (1994) to more than two bidders, and Milgrom and Weber (2000) to imperfect substitutes. We show that this dynamic game has a symmetric pure-strategy Bayes-Nash equilibrium. In this equilibrium, participants bid their value for today's auction less a term that represents their expected surplus from tomorrow's auction—the opportunity cost of winning today. The difficult part of the equilibrium construction (and the reason we need to restrict attention to only two objects throughout the paper) is characterizing the expected surplus function, which is surprisingly subtle in equilibrium. The issue is that *today's* winning bid conveys information about the set of participating bidders in *tomorrow's* auction: since all bidders reduce their bids today as a function of their values of tomorrow's object, losing to a lower bid today makes higher competition tomorrow more likely. At the margin, each bidder therefore needs to assess the opportunity cost of winning today not only as a function of his valuation of tomorrow's object, but also as a function of the bid he submits today. Following Che and Gale (1998), we characterize the equilibrium surplus function in terms of *isobids*—sets of bidder types that submit the same first-round bid. The isobids of our game turn out to be well behaved, facilitating the analysis of comparative statics and the allocative efficiency of the auction. Surprisingly, the equilibrium involves first-stage trade almost surely, a consequence of the informational content of winning the first auction akin to the "loser's curse" (Holt and Sherman 1994, Pesendorfer and Swinkels 1997).

Having characterized the sequential auction with information about future objects revealed, we next turn to results about its efficiency relative to other auction-market designs. Our first efficiency result shows that revealing future objects increases the ex-ante expected efficiency of the sequence of auctions. Specifically, we compare the sequential auction with information about future objects revealed to a sequential auction with information about future objects *hidden*, first analyzed by Engelbrecht-Wiggans (1994). When future objects are hidden, bidders know a second auction will occur and know the distribution of potential valuations in the second auction, but they do not yet know their own specific valuations. The reason why we might expect the information about future objects to enhance efficiency is that informed bidders allocate their demand better across auctions: bidders with high values tomorrow can bid cautiously today, and bidders with low values tomorrow can bid aggressively today. We find that there are always some realizations of bidder values where revealing future objects lowers efficiency, but revealing future objects is indeed good for social welfare when averaged over all possible bidder valuations.[1] Our proof utilizes a classic idea from the theory of single-object auctions with entry, namely, that the expected individual surplus of a participating bidder in a standard auction is equal to his expected contribution to social surplus (McAfee and McMillan, 1987). In our sequential context, self-interested bidders shade their first-round bids precisely by their conditional expectation of second-round individual surplus, which we show is equal to their expected contribution to second-round social surplus. Thus, whenever the winner of the first object is not the bidder who values it most, society is better off in expectation with that "highest valuation" bidder participating in the second auction instead.

Our second efficiency result shows that regardless of information about the second object, sequencing itself increases the expected efficiency of the collection of auctions. We compare the sequential auction to an auction market in which multiple individual auctions are separated, forcing each bidder to choose just one auction to participate in. For instance, imagine that the auctions take place simultaneously in separate rooms. Intuitively, we expect sequencing to be good for efficiency because it allows bidders to participate in more auctions; any bidder who loses the first auction can then participate in the second. As with the case of hidden future objects, it is always possible to find realizations of bidder values where running the auctions simultaneously—hence forcing bidders to choose ex-ante which one auction to participate in—actually increases allocative

[1] It is immediate that revealing private valuations increases efficiency in a standard single-object auction. Moreover, revealing valuations also increases seller revenue as long as there are at least three bidders (Board, 2009).

efficiency. But, in expectation, we show that this kind of marketplace congestion is always bad for welfare.

Our last result shows that the sequential auction with future objects revealed is in a certain sense approximately efficient. More specifically, we show that the expected inefficiency of the sequential auction with future objects revealed is bounded above by the expected efficiency gain associated with adding one more bidder to the second auction. This bound is in the style of Bulow and Klemperer's (1996) bound on the revenue gain from using an optimally set reserve price. An interpretation is that the potential welfare gains from switching to a more sophisticated multi-object auction, such as Vickrey Clarke Groves, are small in this environment.

The remainder of this paper is organized as follows. Section 2 presents the model. Section 3 characterizes the equilibrium of the sequential auction with information about future objects revealed. Section 4 presents our three main results: on information, sequencing, and approximate efficiency. Section 5 concludes.

2 Model

The supply side of the market consists of two objects $j = 1, 2$ for sale. The value the seller of object j derives from keeping her object is normalized to zero. The demand side of the market consists of $N \geq 3$ risk-neutral bidders indexed by i, each with unit capacity. Each bidder i's type is described by a pair of private valuations for the two objects $(x_{i,1}, x_{i,2})$ with the valuations independent across bidders. The unit-capacity constraint makes the two objects substitutes for each bidder; the possibility that $x_{i,1} \neq x_{i,2}$ makes the objects *imperfect* substitutes.[2] The following assumption about the joint distribution of $(x_{i,1}, x_{i,2})$ in the population of bidders is sufficient for all results in this paper to hold:

[2] For instance, the objects could be two different used cars sold on eBay Motors, with $x_{i,1}$ indicating i's value for car 1, $x_{i,2}$ indicating i's value for car 2, and the capacity constraint arising from i only having access to a single parking space

Assumption A1: For each bidder i and object j, valuation $x_{i,j}$ is drawn independently from a continuous distribution $F_j$with a bounded density f_j and full support on a closed interval $[L_j, H_j]$ for some $L_j \geq 0$.

Observe that Assumption A1 allows for the two objects to have different valuation supports, which for instance, could represent a difference in the objects' quality. We need the independence-across-objects part of the assumption to utilize existing results regarding bidding strategies in sequential auctions with future information hidden (Engelbrecht-Wiggans 1994). Our main technical result (characterization of bidding strategies in sequential auctions with future information revealed) relies on a substantially weaker assumption without independence across objects, we leave the exposition of this more general assumption to the relevane subsection of the paper.

Taxonomy of auction markets aggregated from single-object Auctions

An auction market is a set of single-object auctions. For tractability, we focus on markets that sell each object by a second-price sealed-bid auction without reserve.[3]

We consider a two-dimensional taxonomy of auction-market designs based on 1) the timing of the individual auctions (*sequential* or *separate*), and 2) the information available to the bidders about their private valuations of the objects at the beginning of the game (*revealed* or *hidden*). We now define and discuss these two dimensions of our taxonomy.

The first dimension of our taxonomy defines timing of the individual auctions. In *sequential* auctions, j can be interpreted as a time index. All bidders participate first in auction 1, the winner (if any) exits the market, and the remaining bidders then participate in auction 2.[4] In *separate* auctions, each bidder chooses which one auction to participate in; bidders are not allowed / able to participate in both auctions. For instance, imagine that the auctions take place simultaneously

[3] In the context of eBay auctions, this assumption can be motivated by the fact that most bidding on eBay occurs towards the very end of the auction (Roth and Ockenfels, 2001).

[4] Given our unit capacity assumption, the additional assumption that the winner of auction 1 exits the market is most reasonable when disposal is costly or prohibited. For instance, the winner of a procurement auction may be prohibited from transferring the contract, and resale of a used car bought on eBay is costly due to eBay fees and other transactions costs.

in separate rooms.[5] The *separate* condition can be interpreted more broadly as a tractable device for modeling an auction market in which there are obstacles to participating in multiple individual auctions.

The second dimension of our taxonomy defines the information available to bidders at the beginning of the game. When information is *revealed*, each bidder learns his valuation of both objects in the beginning of the game before making any bidding or entry decisions. When information is *hidden*, each bidder learns his valuation for object j only after entering the auction. Specifically: in the sequential auction, bidder i learns his $x_{i,1}$ before submitting his bid in auction 1, and his $x_{i,2}$only after auction 1 concludes but before submitting his bid in auction 2. In separate auctions, bidder i learns his $x_{i,j}$ only after entering auction j. The *hidden* condition can be interpreted more broadly as a modeling convention that captures that auction buyers know they will have future opportunities to trade but are not yet sure of the exact details of these opportunities.

3 Bidding strategies

Table 1 shows the equilibrium bidding strategies $\beta_j$for each market design in the taxonomy, conditioning on bidder i entering auction j. Not all bidders enter all auctions: in sequential auctions, the winner of the first auction does not enter the second auction because of the unit capacity constraint and costly disposal (see footnote 7). In separate auctions, each bidder enters only one auction by construction. The choice-of-auction stage of the separate auction markets with information revealed raises an interesting coordination problem whereby each bidder wants to both enter an auction for which he has a high valuation, but also an auction that other bidders do not want to enter. As with any coordination game, the choice stage thus has multiple equilibria. Our efficiency ranking results will be valid for *any* perfect Bayesian Nash equilibrium of the entry game. Therefore, we do not explicitly analyze the equilibria of the choice-of-auction stage.

[5] One real-world example is Manheim - the dominant U.S. auctioneer of used cars - which conducts auctions in several physical “lanes” simultaneously (Tadelis and Zettelmeyer 2015).

Table 1: Bidding strategies		Bidders' information about their valuations at the beginning of the game	
		Hidden	*Revealed*
Timing of the auctions	*Separate*	$\beta_j\left(x_{i,j}\right) = x_{i,j}$ (Vickrey 1961)	$\beta_j\left(x_{i,j}\right) = x_{i,j}$ (Vickrey 1961)
	Sequential	$\beta_1\left(x_{i,1}\right) = x_{i,1} - \int_{L_2}^{H_2}\int_{L_2}^{x}\left(x-z\right)dF_2^{N-2}\left(z\right)dF_2(x)$ $\beta_2\left(x_{i,2}\right) = x_{i,2}$ (Engelbrecht-Wiggans 1994)	$\beta_1\left(x_{i,1}, x_{i,2}\right) = x_{i,1} - S\left(x_{i,2}, \beta_1\left(x_{i,1}, x_{i,2}\right)\right)$ $\beta_2\left(x_{i,2}\right) = x_{i,2}$ (New result, Theorem 1)

In all but one market design in our taxonomy, the bidding strategies are standard results: the separate auctions are isolated from each other, so bidding one's valuation is a dominant strategy (Vickrey, 1961). For the same reason, the losers of the first auction ($j = 1$) in any sequential setting bid their valuations $x_{i,2}$ in the second auction. The dominant strategy in the second stage means that information disclosure at the end of the first auction does not influence bidding in the second auction.[6] The first-round bidding strategy in sequential markets with information hidden is slightly more involved: since the bidders do not know their own valuations of the second object at the time of bidding in the first auction, they share common knowledge that if they lose the first auction, they will bid their value $x_{i,2} \sim F_2$ in the second auction. Therefore, all bidders expect the same continuation surplus $E_{x_{i,2}}\left[\int_{L_2}^{x_{i,2}}\left(x_{i,2}-z\right)dF_2^{N-2}\left(z\right)\right]$ should they lose the first auction. Engelbrecht-Wiggans (1994) shows that under our assumption A1, it is an iterated conditional dominant strategy equilibrium for bidder i to bid her valuation $x_{i,1}$ net of this expected surplus in the first auction as shown in Table 1. We now turn to characterizing the first-stage bidding strategy in sequential markets with information revealed.

3.1 Bidding in sequential auctions with information about future objects revealed

The received theory of sequential auctions for substitutes focuses either on auctions of several identical units of a good (Milgrom and Weber 2000, Black and de Meza 1992, Katzman 1999), on auctions of heterogeneous goods without information about future goods (Engelbrecht-Wiggans

[6] The second-stage bidding is the same regardless of whether the bidders learn nothing, only the price, or all the bids of the first stage. This is the principal simplification relative to the first-price sealed-bid format, where disclosure of first-stage bids matters (Reiß and Schöndube 2008, Bergemann and Horner 2010).

1994), on the special case with only two bidders and two stochastically equivalent objects (Gale and Hausch 1994), or on vertically differentiated goods (Beggs and Graddy 1997). Our main technical contribution extends the theory to $N \geq 3$ bidders and two *arbitrary* objects. Specifically, our characterization works not only under assumption A1, but even when for each bidder's pair of valuations is drawn independenlty from any joint distribution F with full support and a bounded density on a closed and bounded rectangle [7] Under this more general assumption, we find there always exists a pure-strategy bidding equilibrium with mostly intuitive properties. Throughout the paper, we will use the following extension of notation to indicate the probability mass under a continuous curve in the support of F:

Definition (mass under a curve): For a continuous curve $\psi : [L_2, H_2] \to [L_1, H_1]$, let $F(\psi, z) \equiv \int_{L_2}^{z} \int_{L_1}^{\psi(x_2)} f(x_1, x_2)\, dx_1 dx_2$.

Consider a focal bidder(x_1, x_2), where we suppress the bidder subscript for clarity. We restrict our attention to symmetric equilibria in pure strictly monotone strategies, where strict monotonicity of a bidding strategy is defined as follows:

Definition (strict monotonicity): A bidding strategy $\beta_1(x_1, x_2)$ is *strictly monotone* when it is increasing in x_1 for every fixed $x_2 \in [L_2, H_2]$.

When bidding in the first auction, the focal bidder needs to consider his continuation payoff should he lose the first auction. As in the *Sequential&Hidden* markets discussed above, participation in the second stage yields a non-negative expected continuation surplus to the losers of the first stage. We denote this surplus S. Unlike in the *Sequential&Hidden* markets, the continuation surplus in *Sequential&Revealed* markets depends on both pieces of private information known at the time of bidding. The dependence on x_2 is direct and obvious: higher x_2 implies a higher chance of winning the second auction as well as a higher surplus conditional on winning. The dependence

[7] Under the assumption A1, $F(x_1, x_2) = F_1(x_1) F_2(x_2)$

on $x_1$is indirect via the first-auction bid-level, and it arises because the continuation payoff to losers of the first auction is endogenous to their first-auction bidding strategies in equilibrium: since all bidders reduce their bids today as a function of their values of tomorrow's object, losing to a lower bid today makes higher competition tomorrow more likely. At the margin, each bidder therefore needs to assess the opportunity cost of winning today not only as a function of his valuation of tomorrow's object, but also as a function of the bid he submits today. We now make this intuition about properties of the equilibrium continuation surplus precise by first postulating a particular parametrization of the continuation payoff, and then showing a unique pure-strategy equilibrium with such a structure of the continuation payoff exists.

Definition (first-order regularity): Let $S(x_2, w_1)$ represent the expected continuation surplus to a loser of the first stage who values the second object x_2 and loses the first auction to another bidder's winning bid w_1. We call such a function *first-order regular* when when $S(x_2, L_1) = 0$, S is continuous, and S does not decrease in w_1 weakly faster than unity, i.e. $S(x_2, d) - S(x_2, c) > (-1)(d - c)$ for every $d > c$.

Given any first-order regular $S(x_2, w_1)$, consider a focal bidder(x_1, x_2) who believes the highest bid of his first-auction opponents h_1 is distributed according to some continuous distribution G on $[0, H_1]$, and who also believes his continuation payoff should he lose to a winning bid $w_1$is $S(x_2, w_1)$[8]. Such a bidder solves the following problem in the first auction:

$$\beta_1(x_1, x_2) = \underset{b}{\operatorname{argmax}} \left\{ \int_0^b (x_1 - h_1)\, dG(h_1) + \int_\beta^{H_1} S(x_2, h_1)\, dG(h_1) \right\} \tag{1}$$

When the S function is first-order regular, the following first-order condition characterizes the bidder's best response[9]:

[8]Note that it is not immediately obvious that an *equilibrium* S will satisfy these conditions. For example $S(x_2, L_1) = 0$ seems counter-intuitive at first because a bidder with $x_2 = H_2$ wins the second auction almost surely. We show later that there is a symmetric pure-strategy equilibrium that satisfies all these conditions.

[9] Please see the appendix for detailed analysis of why first-order regularity is sufficient. The argument is straightforward and relies on the Intermediate Value Theorem

$$\beta_1(x_1, x_2) = b \text{ such that } b = x_1 - S(x_2, b) \tag{2}$$

The best response function in equation 2 is intuitive given the truth-revealing property of the second-price auction: the bidder bids her value of the first object net of the opportunity cost of winning, where the opportunity cost of winning the first auction is not being able to participate in the second auction. Equation 2 says that when evaluating the option value of the second auction should he lose, the bidder should bid *as if* he would lose in a tie. In other words, he should assume to be pivotal to the outcome of the first auction. This is the only situation in which changing his first bid slightly changes the outcome of the game, and $S(x_2, b)$ is thus the opportunity cost relevant at the margin.

To see why the magnitude of the first-auction winning bid $w_1$is informative about the continuation payoff *in a symmetric pure-strategy equilibrium*, suppose all bidders bid according to equation 2. Since the competitors' first-auction bids are influenced by their (already revealed) valuations of the second good, knowing that all competitors remaining in the game bid less than w_1 is informative their values for the second item. Therefore, w_1 is informative about the intensity of competition in the second auction. In a Bayes-Nash equilibrium, the beliefs summarized by S must be correct given the first-stage bidding strategyβ_1, and vice versa. The resulting equilibrium restriction is most parsimoniously characterized with *isobids*— sets of valuation pairs that submit the same first-round bid in equilibrium:

Definition (isobid): *Isobid for bid-level* b is a function $I_b(x_2) : [L_2, H_2] \to \mathbb{R}$ such that $\beta_1(I_b(x_2), x_2) = b$ for all $x_2 \in [L_2, H_2]$.

Given the above definition of an isobid, the following proposition describes the equilibrium restriction on first-auction bidding in strictly monotone strategies:

Proposition 1: *For every $b > 0$, an isobid of a symmetric pure-strategy Bayesian Nash equilibrium in strictly monotone strategies must satisfy:*

$$I_b(x_2) = b + \int_{L_2}^{x_2} \left(\frac{F(I_b, z)}{F(I_b, H_2)} \right)^{N-2} dz \tag{3}$$

We include the proof of Proposition 1 in the main body of the paper because the concept of an equilibrium isobid is central to our equilibrium and its properties.

Proof of Proposition 1: Note first that bidding that satisfies equation 2 implies the following form of the isobid for bid-level b:

$$I_b(x_2) = b + S(x_2, b) \tag{4}$$

Consider a focal bidder (x_1, x_2) who loses the first auction to a first-stage winning bid of magnitude $b : w_1 = b$. The winner of the first stage (who submitted the w_1 bid) exits the market, leaving $N - 2$ surviving competitors who also lost the first auction along with the focal bidder. Assume the surviving competitors all follow the same strictly monotone bidding strategy $\beta_1$characterized by equations 2 and 4. Strict monotonicity of bidding strategy β_1 implies that the valuation of each surviving competitor i satisfies $x_{1,i} < I_b(x_{2,i})$; that is, the competitors have valuations $(x_{1,i}, x_{2,i})$ below the isobid I_b. This information can be used to derive the cumulative distribution function of a surviving competitor's valuation of the second object as follows: the probability the surviving competitors i's valuation for the second object $x_{2,i}$ is below some level z is the ratio of the f probability mass below $I_b(x_2)$ and to the left of z, and the entire f probability mass under the isobid $I_b(x_2)$:

$$\Pr(x_{2,i} \leq z \mid w_1 = b) = \frac{\int_{L_2}^{z} \int_{L_1}^{I_b(y_2)} f(y_1, y_2)\, dy_1 dy_2}{\int_{L_2}^{H_2} \int_{L_1}^{I_b(y_2)} f(y_1, y_2)\, dy_1 dy_2} = \frac{F(I_b, z)}{F(I_b, H_2)} \overset{Fig1}{=} \frac{\int_X f(w)\, dw}{\int_{X+Y} f(w)\, dw}, \tag{5}$$

where X and Y are the pertinent areas under the isobid illustrated in Figure 1.

The probability distribution $\Pr(x_{2,i} \leq z \mid w_1 = b)$ and independence across bidders in turn imply the expected continuation surplus $S(x_2, b)$ the focal bidder faces. Following a standard result in auction theory, the expected surplus is simply the integrated probability of winning the second auction:

$$S(x_2, b) = \int_{L_2}^{x_2} (x_2 - z)\, d\mathrm{Pr}^{N-2}(x_{2,i} \leq z \mid w_1 = b) = \int_{L_2}^{x_2} \mathrm{Pr}^{N-2}(x_{2,i} \leq z \mid w_1 = b)\, dx \ , \qquad (6)$$

where the second equality follows from integration by parts. Plugging equation 5 into equation 6 yields the expected continuation surplus of the focal bidder in terms of $I_b(x_2)$, and the result into equation 4 yields the equilibrium restriction in equation 3. *QED Proposition* 1.

The intuition for Proposition 1 is that the isobid for bid-level b implies a belief about second-stage competition, which in turn implies the isobid for the same bid-level b via the best reponses of bidders who assume to be pivotal. In a pure-strategy Bayes-Nash equilibrium, the isobid must be stable given the beliefs if generates. Note that considering isobids instead of directly solving for the bidding function involves a useful dimension reduction: equation 3 shows that an isobid for each bid level depends only on itself and not on the isobids for any other bid levels.

Proposition 1 implies a surprising property of any strictly monotone equilibrium, namely, that $\beta_1(x_1, x_2) \geq L_1$ almost surely:

Corollary to Proposition 1: *In any strictly monotone equilibrium, no bidder abstains from the first auction, and no bidder submits a first-auction bid below* L_1.

We provide both a proof of the Corollary and an intuitive unraveling argument. Suppose a strictly monotone equilibrium exists (which we will establish below under fairly general conditions). Bids below L_1 are ruled out because the dashed line in Figure 1 cannot satisfy the equilibrium condition in Proposition 1: consider a bidder (L_1, z') at the intersection of the dashed isobid curve of some bid-level $b' < L_1$ and the lower boundary of the valuation support: if she is the pivotal bidder

Figure 1: Illustration of an isobid

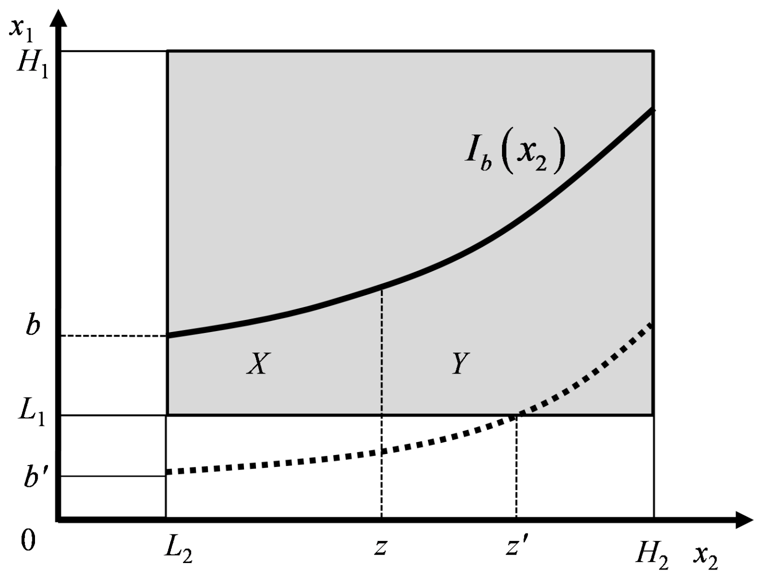

in the first auction, she is guaranteed to lose the second auction by encountering only stronger competitors, so her best response to the dashed curve is to bid her true value $x_1 = L_1$, which contradicts her being on the isobid curve for bid level $b' < L_1$. Thus, symmetric pure-strategy equilibrium considerations alone restrict attention to first-round bids at or above the lower bound of the support of x_1.

If nobody bids below L_1, who are the bidders bidding exactly the lower bound of the support L_1? The equilibrium condition in equation 3 implies that the equilibrium second-stage surplus vanishes as the bid-level approaches L_1, even when $x_2 = H_2$, and so the isobid for bid level L_1 isobid is thus a constant function $I = L_1$. This is counterintuitive because one single-agent best-response thinking would suggest that $x_2 = H_2$has to imply a positive continuation surplus because it guarantees a win of the second auction almost surely. So one would naively expect there to be a bidder with $x_{1,i} > L_1$ and $x_2 = H_2$ on the $I_{L_1}(x_2)$isobid. Such an increasing L_1 candidate isobid unravels as follows: Suppose a focal bidder (x_1, x_2) with $x_1 > L_1$ who sits on the L_1 isobid,

($x_1 = I_{L_1}(x_2) > L_1$), and who is thus conjecturing a positive expected surplus should she lose. In this conjecture, the focal bidder is relying on a positive mass of other bidders $(x_{1,i}, x_{2,i})$ with smaller but still non-minimum valuations $L_1 < x_{1,i} < x_1$ & $L_2 < x_{2,i} < x_2$ to also bid L_1 or less ($L_j < x_{j,i}$ is critical for a positive mass of such competitors). In a symmetric equilibrium, those bidders$(x_{1,i}, x_{2,i})$ rely on other bidders with yet smaller but still non-minimum valuations to also bid L_1, all the way down to bidders arbitrarily near the point (L_1, L_2). But bidders on $I_{L_1}(x_2)$ sufficiently close to (L_1, L_2) realize they will lose the second auction almost surely because the probability mass under $I_{L_1}(x_2)$ and left of a small x_2 is zero (since an equilibrium I must have a slope and curvature of zero at L_2). Therefore, the cascade unravels, the $(x_{1,i}, x_{2,i}) \approx (L_1, L_2)$ competitors bid a small amount strictly greater than L_1, and the original focal bidder thus also bids strictly more than L_1. Any pure abstention strategy unravels in an analogous fashion. See step 2C in the Proof of Theorem 1 for a mathematically rigorous exposition of the unravelling argument.

To close the symmetric equilibrium construction, we need to show that the isobids defined by Proposition 1 exist and imply a well-behaved expected surplus function. This is our main technical result:

Theorem 1: *For any joint distribution $F(x_1, x_2)$ with full support and a bounded density on a closed and bounded rectangle $[L_1, H_1] \times [L_2, H_2] \subset \mathbb{R}_+^2$ and any number $N \geq 3$ bidders drawn independently from F, there is a unique symmetric pure-strategy Bayes-Nash equilibrium in strictly monotone strategies with a continuous bidding function $\beta_1(x_1, x_2)$ that satisfies:*

- $\beta_1(L_1, x_2) = L_1$, $\beta_1(x_1, L_2) = x_1$ *and* $x_1 > \beta_1(x_1, x_2) > L_1$ *for all* $(x_1, x_2) > (L_1, L_2)$,
- $\beta_1(x_1, x_2)$ *is decreasing in* x_2 *for all* $(x_1, x_2) > (L_1, L_2)$

The equilibrium can be characterized by a unique set of equilibrium isobids $\{I_b(x_2)\}_{b=L_1}^{H_1}$, *each of which satisfies equation 3. The equilibrium isobids imply a unique equilibrium expected surplus function* $S(x_2, w_1) = I_{w_1}(x_2) - w_1$, *which is first-order regular. In terms of* S, *the bidding function* $\beta_1(x_1, x_2)$ *satisfies* $\beta_1(x_1, x_2) = x_1 - S(x_2, \beta_1(x_1, x_2))$.

We construct the equilibrium in two steps: First, the full support and boundedness of the joint density imply that for every $b \in (L_1, H_1]$

$, a unique function \mathrm{I} exists that satisfies equation 3. The \mathrm{I} function is a candidate for an isobid curve \mathrm{I}_b (x_2)$

[10] Second, we show that the candidate surplus function $S(x_2, w_1) = I_{w_1}(x_2) - w_1$ implied by the candidate isobids is first-order regular, and so equation 2 thus characterizes the best response to the candidate surplus function. The last part of the proof confirms that the resulting bidding function is increasing in x_1, so the candidate isobids are indeed the equilibrium isobids of a strictly monotone equilibrium.

The most seemingly counter-intuitive property of the equilibrium is highlighted by the above Corollary to Proposition 1: all bidders bid at least L_1, regardless of their x_2. One way to understand the bidding incentives is to consider the informational content of losing the first auction to a low winning bid w_1. Since losing to a lower bid today makes higher competition tomorrow more likely, the first auction involves a "loser's curse" (Holt and Sherman 1994, Pesendorfer and Swinkels 1997) in that a failure to anticipate the informational content of winning makes one bid too low. Specifically, ignoring the information about tomorrow's competition contained in winning with very low bid today would make some bidders bid very low, only to be surprised tomorrow at the intensity of competition.

The first step of the proof (existence and uniqueness of candidate isobids) relies on showing that a K-times repetition of the mapping on the space of functions defined by the RHS of equation 3 is a contraction mapping, so there is a natural numerical method for computing isobids:

Corollary to Theorem 1 : numerical procedure for computing $\beta_1(x_1, x_2)$

The following steps can be used to numerically approximate $\beta_1(x_1, x_2)$ on a grid:

1. Starting with $b = H_1$ and proceeding in small steps of size δ, compute the equilibrium $I_b(x_2)$ for a set of $b \in \{L_1, L_1 + \delta, L_1 + 2\delta, ..., H_1\}$ by iterating equation 3, starting with $I_{b+\delta}(x_2)$.

[10] The fact that the distribution F has a bounded density is important for the existence of candidate isobids; the construction may fail for distributions with atoms or for discrete distributions.

2. Construct equilibrium $S(x_2, w_1)$ by subtracting w_1 from $I_{w_1}(x_2)$.

3. Solve for $\beta_1(x_1, x_2)$ approximately on a grid using equation2.

The first step makes use of continuity of I in b to initialize each set of iterations. Starting with $b = H_1$ is efficient because isobids for high values of b have greater mass of f underneath, so their mapping contracts faster (see proof of Step 1 of Theorem 1).

4 Efficiency bounds and rankings

Having characterized the equilibrium bidding in a *Sequential&Revealed* market, we now turn to our efficiency ranking results. We present two sets of results: First, we show that the *Sequential&Revealed* market generates greater expected surplus than any other market format in our taxonomy. Second, we show that the *Sequential&Revealed* market in a certain sense approximates full efficiency. Both sets of results rely on our workhorse finding that the *Sequential&Revealed* market is more efficient than the greedy allocation, defined as follows:

Definition: The *greedy allocation* assigns the first object to the bidder i with the highest $x_{i,1}$, and assigns the second object to the bidder with the highest $x_{i,2}$ among the remaining bidders.

Theorem 2: *For any joint distribution $F(x_1, x_2)$ with full support and a bounded density on a closed and bounded rectangle $[L_1, H_1] \times [L_2, H_2] \subset \mathbb{R}_+^2$ and any number $N \geq 3$ bidders drawn independently from F, the unique strictly monotone symmetric pure-strategy Bayes-Nash equilibrium of the sequential auction with future objects revealed generates greater expected social surplus than the greedy allocation.*

Proof of Theorem 2: The distributional assumption is needed to guarantee existence of the equilibrium, see Theorem 1. The rest of the proof exploits the properties of equilibrium isobids. Let A be the bidder with the highest value for good 1, B be the bidder with the highest bid in the first auction, and let z be the maximum $x_{i,2}$ among all bidders other than A and B. Let B's bid-level (i.e. the highest bid level) in the first auction be $\bar{b} \equiv \beta_1(x_{B,1}, x_{B,2})$, and let the A's bid-level be $b_A \equiv \beta_1(x_{A,1}, x_{A,2})$. Note that z is the maximum $x_{i,2}$ of $N-2$ bidders, each of whom has $(x_{i,1}, x_{i,2})$ under the isobid for the bid level $\bar{b}$: $\beta_1(x_{i,1}, x_{i,2}) < \beta_1(x_{B,1}, x_{B,2}) \Rightarrow x_{i,1} < I_{\bar{b}}(x_{i,2})$. From equation 6, z is thus exactly the second-auciton competition bidder B expects at the margin should he lose the first auction.

When $A = B$, the two allocations are the same. When $A \neq B$, the greedy allocation generates social surplus of $W_{greedy} = x_{A,1} + \max(z, x_{B,2})$. In contrast, the sequential auction with information revealed generates $W_{rev} = x_{B,1} + \max(z, x_{A,2})$.

We show that for every realization of $(x_{A,1}, x_{A,2})$ and $(x_{B,1}, x_{B,2})$, $E_z(W_{rev} - W_{greedy} | x_{A,1}, x_{A,2}) > 0$. Subtracting z from both surpluses, the difference between the realized social surpluses is

$$W_{rev} - W_{greedy} = \underbrace{x_{B,1} - x_{A,1}}_{loss\,in\,1st\,auction} + \underbrace{\max(0, x_{A,2} - z) - \max(0, x_{B,2} - z)}_{gain\,in\,2nd\,auction}$$

The greedy allocation always realizes more surplus in the first auction. However, the sequential auction realizes more surplus in the second auction as long as $x_{A,2} > x_{B,2}$, which is necessary for $b_{A,1} < b_{B,1}$ and $x_{A,1} > x_{B,1}$ $(A \neq B)$. Since z is the maximum of $N-2$ bidders i with $x_{i,1} < I_{\bar{b}}(x_{i,2})$, taking an expectation over z yields $E_z[\max(0, x - z)] = S(x, \bar{b})$, where S is the equilibrium expected surplus function characterized by Theorem 1. Therefore, the expectation of the difference between surpluses is the (always-positive) difference between the isobids at $x_{A,2}$:

$$E_z(W_{rev} - W_{greedy} | x_{A,1}, x_{A,2}) > 0 = \underbrace{x_{B,1} - x_{A,1}}_{loss\,in\,1st\,auction} + \underbrace{S(x_{A,2}, \bar{b}) - S(x_{B,2} \mid \bar{b})}_{expected\,gain\,in\,2nd\,auction} = I_{\bar{b}}(x_{A,2}) - I_{b_A}(x_{A,2}) > 0$$

where the last equality follows from the definition of an isobid. Please see Figure 4 for a graphical illustration of this proof. *QED Theorem* 2.

Recall that the greedy allocation first sorts the bidders according to their valuations of the first object, so the first auction in the sequence generates the maximum possible social surplus under the greedy allocation. The intuition for Theorem 2 is that whenever the sequential auction with

Figure 2: **Illustration of Theorem 2: Sequential auction with information about future objects revealed is more efficient than the greedy allocation**

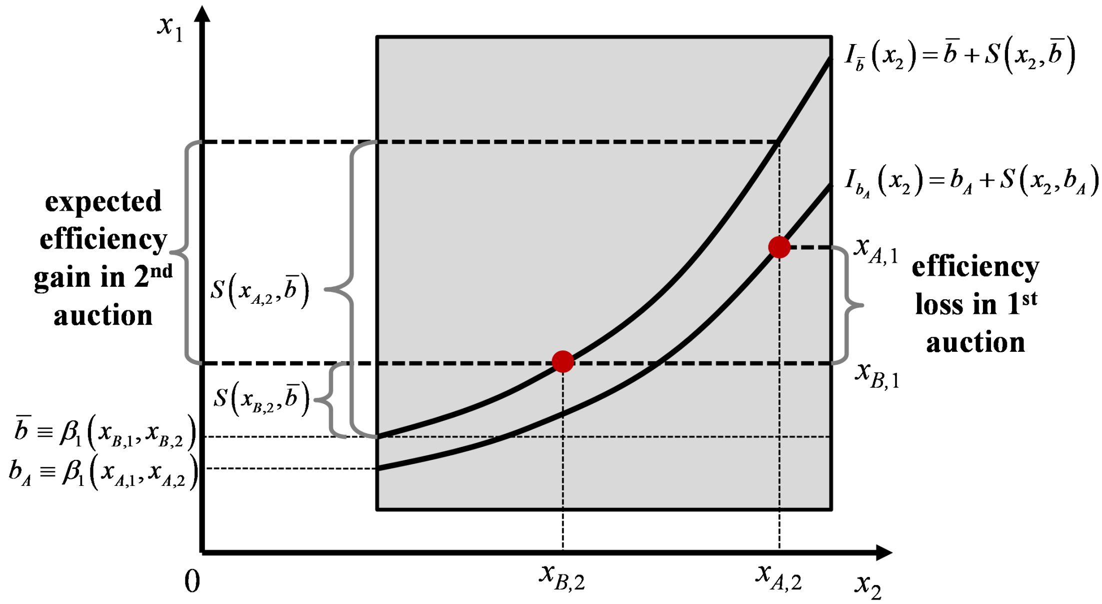

information revealed changes the identity of the first-auction winner (and thus results in a loss in efficiency due to a lower x_1 of the new winner), the second period more than compensates in expectation: the person who would have received the second object under the greedy allocation must have such a high x_2 that his presence in the second auction increases expected social surplus. Note, however, that the efficiency comparison does not hold in *realization*: there always exist profiles of bidders types for which the greedy allocation is efficient whereas the sequential auction with future objects revealed is not (we provide an example of an inefficient realization as part of the overall discussion of our efficiency ranking below).

When each bidder's valuations are independent across the two auctions, the *Sequential&Hidden* auction market results in the greedy allocation. We thus have the following corollary to Theorem 2:

Corollary to Theorem 2: *Under the distributional assumption* A1 *and for any number of bid-*

ders $N \geq 3$, revealing information about future objects increases the expected efficiency of sequential-auction markets.

Theorem 2 implies that the efficiency comparison between *Sequential&Revealed* and *Sequential&Hidden* generalizes at least to any environment in which *Sequential&Hidden* is greedy in equilibrium, i.e. when its bidding strategy $\beta_1(x_1)$ is strictly monotonic, i.e. increasing in x_1. Such a property seems to be a general feature of bidding in the *Sequential&Hidden* market until one considers the impact of possible correlation between $x_{i,1}$and $x_{i,2}$within each bidder i. Such a correlation effectively makes information about future valuations only partially hidden: bidders with different $x_{i,1}$'s will have different beliefs about their $x_{i,2}$'s, and so every bidder i must condition his first-round bid on his belief about $x_{i,2}$, as well as on his belief about other losers' valuations $x_{-i,2}$ on the margin of first-round victory. Thus the elegant analysis of Engelbrecht-Wiggans (1994) no longer goes through when $x_{i,1}$and $x_{i,2}$are correlated. We conjecture that bidding in *Sequential&Hidden* markets is strictly monotonic even when $x_{i,1}$and $x_{i,2}$are arbitrarily correlated and drawn from a joint distribution F that satisfies conditions of Theorem 1, but a formal proof eludes us. Our next result indicates that sequencing itself enhances efficiency:

Theorem 3: *Under the distributional assumption* A1*, the unique strictly monotone symmetric pure-strategy Bayes-Nash equilibrium of the Sequential Auction with information revealed generates greater expected social surplus than does any Perfect Bayes-Nash Equilibrium of the Separate Auctions.*

Proof of Theorem 3: As in the Proof of Theorem 2, let B be the bidder with the highest bid for good 1 in the *Sequential&Revealed* auction, let A be some other bidder, and let , and let z be the maximum x_2 among all bidders other than A and B. As in the Proof of Theorem 2, the expected social surplus in the in the *Sequential&Revealed* auction is:

$$E_z(W_{seq}) = x_{B,1} + E_z(\max(z, x_{A,2})) = x_{B,1} + E_z(\max(0, x_{A,2} - z) + z) = x_{B,1} + S_B(x_{A,2}) + E(z) \tag{7}$$

We now analyze the expected social surplus in the separate auctions. There are two cases, depending on realized entry into the first auction.

Case 1: At least one bidder enters auction 1

If B also wins the first object when the auctions are separated then the separate auctions generate weakly less social surplus because the set of participants in the separated second auction is a subset of the participants in a second auction in a sequence. Since the highest-$x_{i,2}$ participating bidder wins the second object in either equilibrium, separating the second auction must therefore generate weakly less social surplus.

If B does not win the first auction when the auctions are separated, then Let A be the bidder who wins auction 1. The expected social surplus with the auctions separated is bounded by assuming that all bidders other than A enter the second auction: $E_z(W_{sep}) \leq Bound(EW_{sep}) \equiv x_{A,1} + E_z[\max(0, x_{B,2} - z)] + E(z)$. We have shown in the Proof of Theorem 2 that $E_z(W_{seq}) > Bound(EW_{sep})$.

Case 2: No bidders enter auction 1

With all bidders in auction 2, the social surplus is the highest $x_{i,2}$ from all N bidders. Let C be the bidder with highest $x_{i,2}$ other than $x_{B,2}$: $x_{C,2} = \underset{i \neq B}{max} x_{i,2}$.With the auctions sequenced, the social surplus is $W_{seq} = x_{B,1} + x_{C,2}$, whereas the separate auctions yield $W_{sep} = \max(x_{C,2}, x_{B,2})$ in the present case. We show that conditional on any $(x_{B,1}, x_{B,2})$, $E_C(W_{seq}) > E_C(W_{sep})$.

Let $\Phi_{B,2}(x)$ be the *cdf* of a x_2 drawn randomly from below the isobid for bid level $b_{B,1}$. From independence of valuations across bidders, $\Pr(x_{C,2} < z) = \Phi_{B,2}^{N-1}(z)$. The difference between the two expected social surpluses is

$$E_C(W_{seq}) - E_C(W_{sep}) = \underbrace{x_{B,1}}_{sep\, worse\, in\, auction\, 1} - \underbrace{\Pr(x_{C,2} < x_{B,2})\, E_C(x_{B,2} - x_{C,2} \mid x_{C,2} < x_{B,2})}_{sep\, only\, better\, in\, auction\, 2\, when\, x_{C,2} < x_{B,2}} =$$

$$= x_{B,1} - \underbrace{\int_{L_2}^{x_{B,2}} \left(x_{B,2} - x\right) d\Phi_{B,2}^{N-1}\left(x\right)}_{surplus\,with\,additional\,bidder} > x_{B,1} - \underbrace{\int_{L_2}^{x_{B,2}} \left(x_{B,2} - x\right) d\Phi_{B,2}^{N-2}\left(x\right)}_{equilibrium\,surplus} = b_{B,1} \geq L_1 \geq 0$$

The first equality sign follows from the two relative orderings of $x_{B,2}$ and $x_{C,2}$: when $x_{B,2} \leq x_{C,2}$, society is better off by $x_{B,1}$. When, on the other hand, $x_{B,2} > x_{C,2}$, the social surplus difference has an ambiguous sign and amounts to $x_{B,1} - E_C\left(x_{B,2} - x_{C,2} \mid x_{C,2} < x_{B,2}\right)$. The second equality sign highlights that the average $W_{seq} - W_{sep}$ across the two orderings amounts to $x_{B,1}$ minus the difference between $x_{B,2}$ and $x_{C,2}$ whenever $x_{B,2} > x_{C,2}$. The critical insight is that this difference is of the same form as the equilibrium bidding function in equation 2, except the expected surplus part considers $N-1$ rather than $N-2$ opponents (compare with equation 3). Since more opponents imply smaller surplus, the first inequality follows, so the expected $W_{seq} - W_{sep}$ difference exceeds B's first-auction bid. Therefore, $E_C\left(W_{seq}\right) > E_C\left(W_{sep}\right)$. *QED Theorem* 3.

Together with the trivial observation that sequencing increases efficiency when information is hidden, Theorems 2 and 3 indicate that the sequential auction with future objects revealed is the most efficient auction market in our taxonomy:

Summary of efficiency rankings within the taxonomy: *Under the distributional assumption* A1 *and for any number of bidders* $N \geq 3$*, the Sequential&Revealed auction market generates greater expected social surplus than any other auction market in the taxonomy*

Of course, the sequential auction with future objects revealed is not fully efficient. With three bidders and f uniform on the unit square, the worst-case efficiency loss occurs when the bidder values are about (1,0.99), (0,1), and (0.58, 0): the first bidder just loses the first auction to the last bidder, resulting in 0.42 of squandered social surplus. Full efficiency in our setting requires a multi-object auction design, such as the Vickrey-Clarke-Groves auction. Our next result shows, however, that although VCG avoids the worst-case scenario, the *expected* efficiency gain of switching from the sequential auction design to the VCG is in a certain sense small:

Theorem 4: *Under the distributional assumption* A1 *and for any number of bidders* $N \geq 3$,

- *The inefficiency of the Sequential&Revealed market is bounded above by the difference in expected social surplus between an* N *bidder auction and an* $N-1$ *bidder auction for the second object held in isolation.*
- *The Sequential&Revealed market with* $N+1$ *bidders generates greater expected surplus than does the efficient Vickrey-Clarke-Groves mechanism with* N *bidders.*

Proof of Theorem 4: Let $X_j^{(k:n)}$ denote the random variable given by the k-th highest of n draws from F_j. Let W_{VCG}^N denote the expected social surplus in the VCG mechanism with N bidders. W_{VCG}^N is bounded above by the expected social surplus in an economy without the unit-capacity constraint, in which each object is allocated to the bidder who values it the most:
$W_{VCG}^N < E\left(X_1^{(1:n)}\right) + E\left(X_2^{(1:n)}\right)$.

Now let $W_{seq\&rev}^n$ denote the expected social surplus in a sequential auction with information revealed and n bidders. Theorem 2 implies that $W_{seq\&rev}^n$ is bounded below by the greedy allocation that awards the first object to the bidder who values it the most, but then excludes that bidder from the second auction. Under assumption A1, the greedy expected surplus is: $W_{greedy}^n = E\left(X_1^{(1:n)}\right) + E\left(X_2^{(1:n-1)}\right)$, so $W_{seq\&rev}^n > E\left(X_1^{(1:n)}\right) + E\left(X_2^{(1:n-1)}\right)$. Combining the above two bounds with $n = N$ yields $W_{VCG}^N - W_{seq\&rev}^N < E\left(X_2^{(1:N)}\right) - E\left(X_2^{(1:N-1)}\right)$. Combining the above two bounds with $n = N+1$ yields $W_{seq\&rev}^{N+1} - W_{VCG}^N > E\left(X_1^{(1:N+1)}\right) - E\left(X_2^{(1:N)}\right)$. *QED Theorem* 4.

The first part of Theorem 4 says that the inefficiency of the entire two-auction sequence is less than the efficiency loss in the second auction, held in isolation, if $N-1$ rather than N randomly drawn bidders participated in it. This bound becomes tighter as the population variance in x_2 decreases and/or as N increases. Intuitively, the less expected bidder surplus the second auction provides, the more efficient is the sequential auction with information revealed. The second part of Theorem 4 says that a designer concerned about efficiency would rather hold the sequential auction with information revealed than use VCG if the former attracts just a single additional bidder—for example, due to its comparative simplicity. This bound is in the spirit of Bulow and Klemperer (1996), who show in the single-object setting that the revenue benefit of an optimally chosen reserve price is smaller than the revenue benefit of adding one more bidder to the game.

At present we have not been able to generalize Theorem 4 to setting A1. The issue is that when values are correlated we cannot exploit the simple characterization of the greedy expected surplus as $W^n_{greedy} = E\left(X_1^{(1:n)}\right) + E\left(X_2^{(1:n-1)}\right)$. Instead, the best available bound on the greedy surplus is $W^n_{greedy} \geq E\left(X_1^{(1:n)}\right) + E\left(X_2^{(2:n)}\right)$ which in turn implies a weaker inefficiency bound when one follows the proof technique of Theorem 4. Specifically, inefficiency is bounded by the expected difference between the top two valuations of the second object among N bidders.

Simulation results in Budish (2008) suggest that the inefficiency of the sequential auction with future-objects revealed is often substantially smaller than the Theorem 4 bound. For instance, if there are three bidders with values for each object distributed uniformly on [0,1] then VCG achieves expected social surplus of 1.45, whereas the sequential auction with future objects revealed achieves expected social surplus of 1.44.

5 Conclusion

This paper studies how single-object auctions "aggregate up" into a multi-object auction market. Such markets are especially prevalent in settings where there would be large coordination or complexity costs associated with adopting a fully efficient multi-object auction—for example, the costs of coordinating multiple different sellers on eBay or at Sotheby's. We identify two aspects of such auction-market design that unambiguously increase expected marketplace efficiency in our model: First, the individual auctions should be conducted in a sequence; Second, the auctioneer should reveal to the bidders in advance all information about the objects being auctioned. Besides demonstrating that the sequential auction with future objects revealed is more efficient than its alternatives, we also show that its expected efficiency is close to the efficient social surplus. Specifically, the market designer would rather run the sequential auction with future objects revealed than run a fully efficient mechanism but lose one bidder. Together, our results help explain the prevalence of the sequential auction with future objects revealed in practice.

Sequencing auctions and revealing future objects both seem like obvious design decisions in the context of auction markets for substitutes. Both practices have long been standard, for instance, at the classic auction houses Sotheby's and Christie's. Yet many real-life auction markets fail to do one or both of these things. For instance, the now-defunct Amazon and Yahoo auction marketplaces used soft-close ending times, meaning an auction ends only after some time elapses without a new bid (Roth and Ockenfels, 2001). This design made it difficult for bidders to guess which of two auctions would end first, and hence difficult to participate in both. Google Base's auctions—also defunct—were sorted by search-term closeness-of-fit rather than by ending time, making identifying, let alone participating in, the full sequence difficult for bidders. Moving against our efficiency result, in 2008 eBay switched from sorting auctions strictly by ending time ("Ending Soonest") to a sort order based on a variety of features ("Best Match"); interestingly, eBay appears to have at least partly switched back in some product categories.[11] Large wholesale car auctions in the United States often conduct several auctions concurrently, each in a separate "lane" (Tadelis and Zettlemeyer, 2015). Perhaps the starkest example of a non-sequential auction is that run by the event-ticket marketplace StubHub in 2006. The auction was for all of the tickets to a single event, with the tickets sold in pairs.Its single-object auction design was similar to eBay's, but the individual auctions all had *identical* hard-close ending times, and as a consequence its performance was poor.[12] Auctions at charity benefits are often organized similarly to StubHub's, which might explain why the final moments of so-called "silent auctions" are often anything but.

Perhaps the starkest example of a sequential auction with future objects hidden is that described by Engelbrecht-Wiggans (1994), in which future items in an equipment auction are hidden behind

[11] We searched for event tickets and cars on eBay in June 2011. In the event tickets category (e.g., search term "U2 Tickets"), the default "Best Match" display sorted upcoming auction listings in strict order of ending time. In the automobiles category, the default "Best Match" display seemed to place little weight on auction ending time. For instance, for the search terms "Ford F150" and "Toyota Corolla", we found that in 78 percent of our searches (conducted once per day for one week), at least one of the three auctions ending soonest was not shown in the top 25 "Best Match" listings. In 42 percent of the searches, none of the three was.

[12] 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 prices 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).

a curtain until their turn for sale. The field experiment reported by Tadelis and Zettlemeyer (2015) relates to improving the disclosure of information about objects in the sequence, specifically releasing more information about automobiles' quality condition. Some gimmicky online auction sites with very short auction durations (e.g., Bidz.com) seem to purposefully suppress information about future objects for sale.[13] One could also interpret the non-sequential sort used by Google Base and eBay's "Best Match" as obscuring information about the full set of objects for sale.

A limitation of our analysis is that it focuses on the case of just two auctions and a fixed set of bidders. It would be desirable to allow for a longer (perhaps infinite) sequence of auctions, and to allow bidders to arrive and depart from the market at different times. Unfortunately, such analysis quickly becomes intractable due to the asymmetry and learning issues first raised in the seminal paper of Milgrom and Weber (2000);[14] as is clear from our characterization in Section 3, even the case of two auctions and a fixed set of bidders is quite involved. Our intuition is that our efficiency results would be robust to longer sequences, but this remains conjecture.

We close by noting an interpretation of our analysis that suggests a fruitful direction for future research. In the context of online marketplaces (cf. Levin, 2011), sequencing auctions and revealing future object information can be interpreted as aspects of what e-commerce firms call *user-interface design*. Surely there are other elements of user-interface design that have an important impact on market performance: helping users search for objects efficiently; helping sellers describe their objects effectively; reducing the transactions costs, for example, measured in user time, associated with performing various actions in the marketplace; and so on. Our paper perhaps represents a first foray into the economic analysis of this important aspect of online market design.

[13]On Bidz.com, "auctions start at $1 every 5 seconds!" The company has been accused of fraud and faced numerous other difficulties (Miniter, 2008).

[14]Suppose we have three auctions instead of only two, with each bidder having private single-item valuations (x_1, x_2, v_3) (x_1, x_2, v_3), and suppose further that only the winner and the price p_1 of the first unit are revealed before the second stage. Then second-stage bidders have asymmetric beliefs about the v_3 of the remaining competitors, because one of the remaining bidders bid exactly p_1, whereas the other bidders bid strictly below p_1. Even when this asymmetry is resolved by revealing all first-stage bids, as suggested by Milgrom and Weber (2000), the information about future goods would make it necessary to explicitly model second-stage beliefs about v_3: first-stage bids would be a function of v_3, and bidders thus may have an incentive to mislead their competitors into thinking their v_3 is very high by bidding very low in the first stage. Furthermore, the price p_1 would enter second-stage bids, so the last two auctions do not reduce to the environment studied here. Thus, even with just three auctions it is unclear whether there exist intuitive pure-strategy equilibria.

Appendix

Claim: When S is first-order regular, the FOC condition in equation 2 characterizes the best response function to S.

Proof: To see that the $\beta_1(x_1, x_2)$ is well defined by the FOC, note that $S(x_2, d) - S(x_2, c) > (-1)(d - c)$ implies the RHS of equation 2 does not increase in β weakly faster than unity. $S(x_2, L_1) = 0$ fixes the intercept of the RHS at $x_1 - S(x_2, L_1) = x_1 \geq L_1$. Since the LHS of equation 2 is an identity function, continuity of S in b implies (via the Intermediate Value Theorem) that the RHS must eventually intersect the LHS at some $b \geq L_1$ that solves equation 2. Such a solution is unique because the slope of $x_1 - S(x_2, b)$ in b is less than unity everywhere, so a second intersection is impossible. Please see Figure below for an illustration: The solid diagonal line is the LHS of equation 2. The two dashed lines illustrate two possible RHS of equation 2: the lower dashed line corresponds to S increasing in bid-level; the upper dashed line corresponds to S locally decreasing in bid-level, but slower than unity.

Figure 3: The bidding function is well defined by the first-order condition

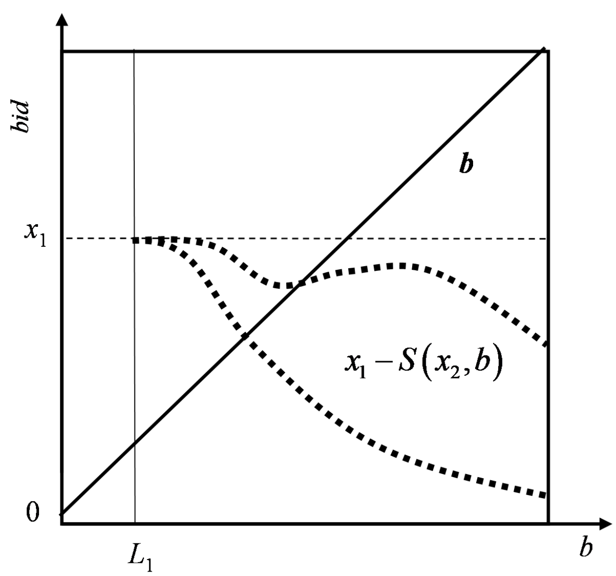

It remains to be shown that $\beta_1(x_1, x_2)$ is the best response to S. The objective function shown in equation 1, which we will denote as $\Pi(b \mid x_1, x_2, S)$, is concave at the FOC when S does not decrease in b faster than unity. To see this, let b_{FOC} be the value of b implied by the FOC. When S is partially differentiable, the SOC is $\frac{d^2\Pi}{db^2} \mid_{b=b_{FOC}} = -1 - \frac{\partial S}{\partial b} \mid_{b=b_{FOC}} < 0 \Longleftrightarrow \frac{\partial S}{\partial b} \mid_{b=b_{FOC}} > -1$. When S is not partially differentiable, $S(x_2, d) - S(x_2, c) > (-1)(d-c)$ is sufficient for concavity of $\Pi(b \mid x_1, x_2, S)$ at the FOC. However, $\Pi(b \mid x_1, x_2, S)$ may not be globally concave. We have shown above that the solution to the FOC is unique, so the only candidates other than b_{FOC} for a true global best response are $b = 0$ and $b = x_1$. It is obvious $\Pi(b_{FOC} \mid x_1, x_2, S) > \Pi(x_1 \mid x_1, x_2, S)$. The single crossing property outlined above also shows that $b = 0$ or even abstaining from the first auction cannot be a profitable deviations from b_{FOC} either. Suppose $N-1$ competitors bid according to β_1, and a bidder i deviates to an abstention *instead of* bidding $B_i \equiv \beta_1(x_{1,i}, x_{2,i})$. The only way such a deviation can make any difference in the outcome of the game is if the focal bidder i would have won the first auction, that is, if $B_i > h_1$ where h_1 is the highest opponent bid as in equation 1. For $B_i > h_1$, the deviation thus yields the expected second-stage surplus of $\int_0^{B_i} S(x_{2,i}, h_1)\, dG(h_1)$ instead of $\int_0^{B_i} (x_{1,i} - h_1)\, dG(h_1)$ available from bidding B_i. But the profit from the deviation must be lower, because $x_{1,i} - b > S(x_{2,i}, b)$ for all $b < B_i$. *QED Claim.*

Proof of Theorem 1

We construct the equilibrium in two steps:

1. The full support and boundedness of f imply that for every $b \in (L_1, H_1]$

$, a unique function \mathrm{I}_b$ exists that satisfies equation 3. This function is a candidate for an isobid curve $I_b(x_2)$.

The set of candidate isobids $I_b(x_2)$ for all $b \in (L_1, H_1]$ implies a candidate surplus function $S(x_2, w_1) = I_{w_1}(x_2) - w_1$. The full support and boundedness of the joint density f, together with the properties of the candidate isobids implied by equation 3 ensure that the candidate surplus function is first-order regular, so there is a unique best response function defined implicitly by $\beta_1(x_1, x_2) = x_1 - S(x_2, \beta_1(x_1, x_2))$, which is strictly monotone.

It is easy to show that the bidding function is increasing in x_1,so the candidate isobids from Step 1 are indeed equilibrium isobids of a pure strategy Bayes-Nash equilibrium in strictly monotone strategies.

Step 1: Existence and uniqueness of candidate isobids

Claim 1: *For each* $b \in (L_1, H_1]$, *there exists a unique nondecreasing* 1*-Lipschitz function* $I_b(x_2) : [L_2, H_2] \to [b, b + H_2 - L_2]$ *with* $I(L_2 \mid b) = b$ *such that* $I_b = \boldsymbol{T}(I_b; b)$ *where* $\boldsymbol{T}$ *is a function in the space of functions parametrized by* b*, and defined by the RHS of equation 3:* $\boldsymbol{T}(J; b)(x_2) = b + \int_{L_2}^{x_2} \left(\frac{F(J,z)}{F(J,H_2)}\right)^{N-2} dz$

Proof: Fix any $b \in (L_1, H_1)$ and denote the global bound on f by $\overline{f}$. Let Ψ be a set of nondecreasing 1-Lipschitz functions $I : [L_2, H_2] \to [b, b + H_2 - L_2]$ such that $I(L_2) = b$: $\Psi = \{I : [L_2, H_2] \to [b, b + H_2 - L_2], I(L_2) = b, \forall x > y, 0 \leq I(x) - I(y) \leq x - y\}$. Ψ is a closed subset of the complete metric space of all bounded continuous functions from $[L_2, H_2]$ to $[0, H_1 + H_2 - L_2]$ with the supremum metric $d(I, J) = \max\limits_{x \in [L_2, H_2]} |I(x) - J(x)|$, so it is itself a complete metric space with the same metric. It is immediate from equation 3 that $\boldsymbol{T}$ projects Ψ into itself: $\boldsymbol{T} : \Psi \to \Psi$.

The remainder of the proof shows that for every $b \in (L_1, H_1)$, a $K \geq 1$ exists such that $\boldsymbol{T}^K(I; b)$ is a contraction map. That is, a $q < 1$ exists such that $d\left(\boldsymbol{T}^K(I; b), \boldsymbol{T}^K(J; b)\right) < qd(I, J)$ for all $I, J \in \Psi$. By the Banach Fixed Point Theorem, this is enough to show that $\boldsymbol{T}$ has a unique fixed point in Ψ, and iterations of $\boldsymbol{T}$ starting at any point in Ψ converge to the unique fixed point exponentially fast. The unique fixed point is the unique candidate isobid $I_b(x_2)$. It is enough to consider $N = 3$, because $d(\boldsymbol{T}(I; b), \boldsymbol{T}(J; b))$ decreases in N (exponentiating the probability in $\int_{L_2}^{x_2} Pr^{N-2}(...)\,dz$ by the $(N-2)$ reduces $\boldsymbol{T}(I; b)$ as N increases and the same is true for $\boldsymbol{T}(J; b)$, so $d(\boldsymbol{T}(I; b), \boldsymbol{T}(J; b))$ decreases in N).

Pick any distance $\delta > 0$ and any $I \in \Psi$, and consider the following bound on how far apart $\boldsymbol{T}(I)$ and $\boldsymbol{T}(J)$ can be pointwise:

$$|\boldsymbol{T}(I; b)(x_2) - \boldsymbol{T}(J; b)(x_2)| = \left| \int_{L_2}^{x_2} \left[\frac{F(I, z)}{F(I, H_2)} - \frac{F(J, z)}{F(J, H_2)} \right] dz \right|$$

$$\overset{(I>b)}{<} F_1^{-1}(b) \left| \int_{L_2}^{x_2} \int_{L_2}^{z} \int_{J(y_2)}^{I(y_2)} f(y_1, y_2)\, dy_1 dy_2 dz \right| \overset{(f<\overline{f})\&(L_2 \geq 0)}{<} \delta \overline{f} F_1^{-1}(b) \int_0^{x_2} z dz = \frac{\delta \overline{f} x_2^2}{2F_1(b)} \qquad (8)$$

The intuition for the bound arises from the worst-possible scenario whereby the denominators are as small as possible, the two curves I and J are δ apart, and there also happens to be a lot of mass between the two curves, and this mass is "integrated over" twice by definition of the $\boldsymbol{T}$ mapping. Intuitively, the two images thus have to be very close together for small x_2 and can only diverge from each other monotonically as x_2 increases. It is immediate that $\boldsymbol{T}$ alone may not be a contraction map based solely on the bound in equation 8, because the equation only shows that $d\left(\boldsymbol{T}\left(I;b\right),\boldsymbol{T}\left(J;b\right)\right) < \frac{\delta\overline{f}H_2^2}{2F_1(b)}$, and $\frac{\overline{f}x_2^2}{2F_1(b)}$ may not be lower than unity. However, note that $\boldsymbol{T}$ bounds the difference between the images of I and J quadratically as a function of x_2. Therefore, iterating the $\boldsymbol{T}$ mapping twice (which we denote by a "2" superscript) results in a quartic bound:J;b

$$\left|\boldsymbol{T}^{\boldsymbol{2}}\left(I;b\right)\left(x_2\right)-\boldsymbol{T}^{2}\left(J;b\right)\left(x_2\right)\right| = \left|\int\limits_{L_2}^{x_2}\left[\frac{F\left(\boldsymbol{T}\left(I;b\right),z\right)}{F\left(\boldsymbol{T}\left(I;b\right),H_2\right)}-\frac{F\left(\boldsymbol{T}\left(J;b\right),z\right)}{F\left(\boldsymbol{T}\left(J;b\right),H_2\right)}\right]dz\right|$$

$$\overset{(I>b)}{<} F_1^{-1}(b)\left|\int\limits_{L_2}^{x_2}\int\limits_{L_2}^{z}\int\limits_{T(J;b)(y_2)}^{T(I;b)(y_2)} f\left(y_1,y_2\right)dy_1dy_2dz\right| \overset{(f<\overline{f})\&(L_2\geq 0)\&eq.8}{<} \overline{f}F_1^{-1}\left(b\right)\int\limits_{L_2}^{x_2}\int\limits_{L_2}^{z}\frac{\delta\overline{f}y_2^2}{2F_1\left(b\right)}dy_2dz = \frac{\delta\overline{f}^2x_2^4}{4!F_1\left(b\right)} \tag{9}$$

By induction, we can thus show that $\left|\boldsymbol{T}^{\boldsymbol{K}}\left(I;b\right)\left(x_2\right)-\boldsymbol{T}^{K}\left(J;b\right)\left(x_2\right)\right| < \delta\left(\frac{\overline{f}x_2^2}{F_1(b)}\right)^K\frac{1}{(2K)!}$.

Since $\lim\limits_{K\to\infty}\frac{C^K}{(2K)!}=0$ for every positive constant C, a K exists high enough that $\boldsymbol{T}^K$ is a contraction, namely K such that $\left(\frac{\overline{f}H_2^2}{F_1(b)}\right)^K\frac{1}{(2K)!}<1$. One way to picture the contraction is that repeated iteration of $\boldsymbol{T}$ brings the images of any two functions in Ψtogether like a zipper closing from left to right along the $x_2$axis. *QED Step 1.*

Step 2: First-order regularity of the candidate expected surplus function implied by candidate isobids

The full set of candidate isobids $\left\{I_b\left(x_2\right)\right\}_{b>L_1}^{H_1}$ from Step 1 implies a unique candidate for the expected surplus function defined on $[L_2,H_2]\times(L_1,H_1]$: $S\left(x_2,w_1\right)=I_{w_1}\left(x_2\right)-w_1$. Extend the surplus function to the entire closed support rectangle by defining it as a limit $S\left(x_2,L_1\right)=$

$\lim_{w_1 \to L_1^+} I_{w_1}(x_2) - L_1$. The resulting S has the following properties stemming from the fact that S is an integrated cumulative distribution function: $S(x_2, w_1) \geq 0$, S is twice partially differentiable in x_2, S is non-decreasing in x_2, 1-Lipschitz in x_2, and convex in x_2:

$S(x_2, w_1) = \int_{L_2}^{x_2} Pr^{N-2}[y_2 < z | y_1 < I_b(y_2)]\, dz \Rightarrow \frac{\partial S}{\partial x_2} = Pr^{N-2}[y_2 < x_2 | y_1 < I_b(y_2)] \geq 0$ and $\frac{\partial^2 S}{\partial x_2^2} = (N-2)\, Pr^{N-3}[y_2 < x_2 | y_1 < I_b(y_2)] \frac{\int_{L_1}^{I_b(x_2)} f(y_1, x_2) dy_1}{F(I_b, H_2)} \geq 0$

where all inequalities are strict when $x_2 > L_2$. The intuition is straightforward: the expected surplus increases in the valuation x_2 because a higher valuation makes winning more likely and also increases the actual surplus conditional on winning. Since these two at-least-linearly increasing components effectively multiply to produce the expected surplus, the convexity results. Since increasing x_2 by a small amount can increase the expected future surplus at most by that amount (and that only in the case when future prices are guaranteed to be below x_2), the slope of S in x_2 is bounded above by unity. The rest of Step 2 of the proof shows that S is well behaved not only as a function of x_2, but also as a function of w_1, namely that it is first-order regular (see Definition of first-order regularity in the main body of the paper):

Claim 2A: $\forall c < d \in (L_1, H_1)\, \& \forall x_2 \in [L_2, H_2]$, $S(x_2, d) - S(x_2, c) > -(d - c)$.

To show that S does not decrease in w_1 weakly faster than unity, it is enough to show that two candidates for equilibrium isobids cannot cross or even touch each other, that is, that $I(x_2, w_1)$ is strictly increasing in both arguments: $\forall c < d \in (L_1, H_1)\, \& \forall x_2 \in [L_2, H_2]$: $d > c \Rightarrow I_d(x_2) > I_c(x_2)$.

Proof: Suppose there is a pair $d > c$ and an $x_2 \in [L_2, H_2]$ such that $I_d(x_2) \leq I_c(x_2)$. Continuity of isobids implies that the two isobids must intersect at or below x_2. Let the smallest intersection point be $x*$, namely, $I_d(x_2) > I_c(x_2)$ for all $x < x*$ and $I_d(x*) = I_c(x*)$. Therefore, I_c intersects I_d from below at $x*$, so the slope of I_c at $x*$ must be weakly higher than the slope of I_d at $x*$:

$$I_c^{'}(x*) \geq I_d^{'}(x*) \tag{10}$$

. The equilibrium restriction (equation is inequality based on the equilibrium relation 3) and full support of f rule this ordering of slopes at $x*$ out. Three distinct cases exist:

Case 1 (single intersection): $I_d \leq I_c$ on $(x*, H_2]$ (see dotted line in Figure accompanying this proof). From equation 3, the slope of an isobid is a probability: $I_b^{'}(x*) = \left(\frac{F(I_b, x*)}{F(I_b, H_2)}\right)^{N-2}$. Let $A = F(\min(I_c, I_d), H_2)$, $A1 = F(I_c, x*)$, and $X = F(I_d, x*) - F(I_c, x*)$ (see Figure 2 for an illustration of these probability masses). From full support of f, $X > 0$. Since $I_d \leq I_c$ on $(x*, H_2]$,

$$I_d^{'}(x*) = \left(\frac{A1 + X}{A + X}\right)^{N-2} > \left(\frac{A1}{A}\right)^{N-2} \geq \left(\frac{A1}{A + B + Y}\right)^{N-2} = I_c^{'}(x*) \tag{11}$$

, where $B + Y \geq 0$ is the probability mass between the two isobids on $(x*, H_2]$, and it is non-negative because of full support and $I_d \leq I_c$ on $(x*, H_2]$. The inequality 11 is a contradiction with the ordering of slopes necessary for an intersection at $x*$as derived in inequality 10. Intuitively, a single intersection at $x*$ forces the conditional probability of a $x_2 < x*$ greater for the higher and flatter isobid, and this conditional probability happens to be the slope of the isobid at $x*$.

Figure 4: Illustration of proof that isobids cannot intersect (Cases 1 and 2 of Step 2A)

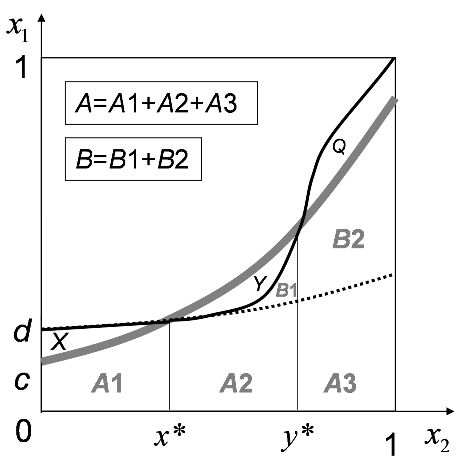

Case 2 (multiple crossings): Suppose the ordering of the isobids on $(x*, H_2]$ is ambiguous. At least one more intersection must be at point $y* > x*$such that $I_c$is crossing $I_d$from above and so

$$I_c^{'}(y*) \leq I_d^{'}(y*) \tag{12}$$

. For equilibrium isobids, the order of slopes at the lowest intersection $x*$shown in equation 10 and the full support rule out this ordering of slopes at $y*$. Let $A2 + B1 = F(I_d, y*) - F(I_d, x*)$, $Y = [F(I_c, y*) - F(I_c, x*)] - [F(I_d, y*) - F(I_d, x*)]$, and $Q = [F(I_d, H_2) - F(I_d, y*)] - [F(I_c, H_2) - F(I_c, y*)]$ (see Figure). From full support of f, $A2 + B1 > 0$ and $Y > 0$, but the sign of Q is ambiguous. Nevertheless, we can express $I_c^{'}(x*) \geq I_d^{'}(x*)$ as

$$I_c^{'}(x*) = \left(\frac{A1}{A+B+Y}\right)^{N-2} > \left(\frac{A1+X}{A+B+X+Q}\right)^{N-2} = I_d^{'}(x*) \tag{13}$$

To obtain the implied slopes at $y*$, add $A2+B1+Y$ to the LHS of equation 13 numerator and only $A2 + B1$ to the RHS of equation 13 numerator. These additions clearly preserve the inequality in equation 13, but they imply an ordering of slopes at $y*$:

$$I_c^{'}(y*) = \left(\frac{A1+A2+B1+Y}{A+B+Y}\right)^{N-2} > \left(\frac{A1+X+A2+B1}{A+B+X+Q}\right)^{N-2} = I_d^{'}(y*) \tag{14}$$

, and this ordering is a contradiction with the necessary ordering of slopes for a second intersection at $y*$ in equation 12.

Case 3: The only remaining possibility is that $x*$ is in fact a point of tangency, namely, $I_c^{'}(x*) = I_d^{'}(x*)$ and $I_d > I_c$ everywhere other than at $x*$. Since Cases 1 and 2 rule out any intersections, tangency at $x*$ means the tangency holds for all $z \in [c, d]$: $I_z^{'}(x*) = \lambda$ for some constant λ. Recall that the candidate isobids are convex in x_2. Since higher z have higher intercepts and there are no other intersections, the curvatures of all the I_z at $x*$ thus must be nondecreasing in z: $\frac{dI_z^{''}(x*)}{dz} \geq 0$. Once again, equilibrium implies this ordering of curvatures cannot happen: for every $z \in [c, d]$, $I_z^{'}(x*) = \left(\frac{F(I_z, x*)}{F(I_z, H_2)}\right)^{N-2} = \lambda$, so

$$I_z^{''}(x*) = (N-2)\,\lambda \left(\frac{\int_{L_1}^{I_z(x*)} f(x_1, x*)\,dx_1}{F(I_z, x*)}\right). \tag{15}$$

Since $I_z(x*)$ is a constant for $z \in [c, d]$ and $F(I_z, x*)$ strictly increases in z because of full support, the numerator of the ratio in equation 15 is constant in z while the denominator increases, and so

$I_z^{''}(x*)$ strictly decreases in z, a contradiction with $\frac{dI_z^{''}(x*)}{dz} \geq 0$. *QED Claim 2A.*

Claim 2B: S is continuous in w_1 on $w_1 \in [L_1, H_1]$.

We proceed in two sub-steps: first, we show that the candidate equilibrium surplus function $S(x_2, w_1)$ is continuous in w_1 at all $w_1 \in (L_1, H_1]$

$.By defining$ $\mathrm{I}_{L_1}(x_2)$ as a limit, we then extend the candidate surplus function to be continuous on the entire closed support $[L_2, H_2] \times [L_1, H_1]$.

To show that $S(x_2, w_1)$ is continuous in w_1 at all $w_1 \in (L_1, H_1]$

$, it is enough to show that$ $\mathrm{I}(x_2|w)$ is upper semi-continuous in $w_1$at all $w_1 \in (L_1, H_1]$

$.The proof of lower semi-continuity is analogous. Fix$ w_1 and consider any monotonically decreasing sequence of $\delta_n > 0$ such that $\lim_{n\to\infty} \delta_n = 0$. The corresponding sequence of isobids $\{I_{w_1+\delta_n}(x_2)\}_{n=1}^{\infty}$ is uniformly bounded because it projects into a closed interval, and it is equicontinuous because all isobids are nondecreasing and have slopes less than unity (1-Lipschitz). Therefore, the Arzela-Ascoli Theorem implies that $\{I_{w_1+\delta_n}(x_2)\}_{n=1}^{\infty}$ has a uniformly convergent subsequence, and its limit is some 1-Lipschitz function $I_{w_1}^{+}(x_2)$. Monotonicity of the original sequence (from candidate isobids not intersecting, shown above in Claim 2A) implies that the subsequence also converges uniformly to $I_{w_1}^{+}(x_2)$, because when $I_{w_1+\delta_n}(x_2)$ is an element of the convergent subsequence, all $I_{w_1+\delta_{n+k}}^{+}(x_2)$ are between $I_{w_1+\delta_n}(x_2)$ and $I_{w_1}^{+}(x_2)$ in the supremum metric. Monotonicity further implies that for every δ, $I_{w_1}(x_2) \leq I_{w_1}^{+}(x_2) < I_{w_1+\delta_n}(x_2)$. For upper semi-continuity in w_1, it remains to be shown that the first inequality is in fact an equality: $I_{w_1}(x_2) = I_{w_1}^{+}(x_2)$.

The uniform convergence of $\{I_{w_1+\delta_n}(x_2)\}_{n=1}^{\infty}$ implies that the implied sequence of distributions $\{\Pr(x_2 \leq z \mid x_1 < I_{w_1+\delta_n}(x_2))\}_{n=1}^{\infty}$ also approaches the implied $\Pr\left(x_2 \leq z \mid x_1 < I_{w_1}^{+}(x_2)\right)$ uniformly. The uniform convergence of $\{\Pr(x_2 \leq z \mid x_1 < I_{w_1+\delta_n}(x_2))\}_{n=1}^{\infty}$ in turn implies that the equilibrium relation from equation 3, which holds for every member of the sequence, is preserved in the limit $I_{w_1}^{+}(x_2) = \boldsymbol{T}\left(I_{w_1}^{+}(x_2); w_1\right)$, and so it must be true that $I_{w_1}^{+}(x_2) = I_{w_1}(x_2)$, because each isobid is unique (Claim 1). *QED Claim 2B.*

Claim 2C: $\lim_{w_1 \to L_{1+}} S(x_2, w_1) = 0$.

The equilibrium condition in equation 3 is not defined for the function $I_{L_1}(x_2) = L_1$, because the probability mass below it is zero. Recall the definition of a cumulative distribution function under a continuous curve from equation **??**. The equilibrium restriction can be equivalently expressed (by differentiating both sides twice) as a differential-integral equation:

$$I_b^{''}(x_2) F^{N-2}(I_b, H_2) = (N-2) F^{N-2}(I_b, x_2) \int_{L_1}^{I_b(x_2)} f(y_1, y_2)\, dy_1 \tag{16}$$

with initial conditions $I_b(L_2) = b$, $I_b^{'}(L_2) = 0$. The alternative equilibrium condition 16 holds for$I_{L_1}(x_2) = L_1$, so it is a candidate for an equilibrium isobid.

We now show that it is a unique candidate, and that higher $I_b(x_2)$ converge to it as b approaches L_1. By the same arguments used above in the proof of Claim 2B, any monotonic sequence $\{I_{L_1+\delta_n}(x_2)\}_{n=1}^{\infty}$ converges uniformly to some 1-Lipschitz function $J(x_2)$ such that $L_1 \leq J(x_2) < I_{L_1+\delta_n}(x_2)$ and $J(L_2) = L_1$. It remains to be shown that $J(x_2) = I_{L_1}(x_2) = L_1$. Suppose $J(x_2) > L_1$ for some $x_2 > L_2$, so there is a positive mass under $J : F(J, H_2) > 0$ (from continuity of J together with full support of f). $F(J, H_2) > 0$ in turn implies that the RHS equation 3 is well defined, so $J = \boldsymbol{T}(J; L_1)$; that is, J is a valid isobid for the bid-level L_1 (again, see proof of Claim 2B for details). The rest of this proof shows no $J : [L_2, H_2] \to [L_1, L_1 + H_2 - L_2]$ can exist such that $F(J, H_2) > 0$, and $J = \boldsymbol{T}(J; L_1)$. Suppose otherwise, let$\delta = d(J, L_1)$ under the supremum metric, and consider the following bound on how far apart L_1 and $\boldsymbol{T}(J; L_1)$ can be pointwise:

$$|L_1 - \boldsymbol{T}(J; L_1)(x_2)| = \left| \int_{L_2}^{x_2} \left[\frac{F(J, z)}{F(J, H_2)} \right] dz \right| < \left| \int_{L_2}^{x_2} \left[\frac{\int_{L_2}^{z} \int_{L_1}^{L_1+\delta} f(y_1, y_2)\, dy_1 dy_2}{F(J, H_2)} \right] dz \right|$$

$$\overset{(f<\overline{f})\&(L_2 \geq 0)}{<} \delta \overline{f} F^{-1}(J, H_2) \int_0^{x_2} z dz = \frac{\delta \overline{f} x_2^2}{2F(J, H_2)} \tag{17}$$

Analogously with equation 8 bounding the images of two arbitrary isobids of the same level to be quadratically far apart, equation 17 bounds the distance between L_1 and $\boldsymbol{T}(J; L_1)$ to be quadratically far apart. We can thus use the same "trick" to bound the distance between L_1 and

$\boldsymbol{T^k}(J; L_1)$, and find that the distance $d\left(\boldsymbol{T^k}(J; L_1), L_1\right)$ must approach zero. The only subtle difference is that L_1 is not a fixed point of $\boldsymbol{T}(\cdot; L_1)$ because $\boldsymbol{T}(L_1; L_1)$ is not defined: $\boldsymbol{T}(\cdot; L_1)$ is not a continuous function on the space of 1-Lipschitz functions $J : [L_2, H_2] \to [L_1, L_1 + H_2 - L_2]$. *QED Claim 2C.*

Step 3: Claims 1 and 2 show that there is a unique first-order regular $S(x_2, w_1)$ on $[L_2, H_2] \times [L_1, H_1]$, so equation 2 characterizes the best response to S that satisfies $\beta_1(x_1, x_2) = x_1 - S(x_2, \beta_1(x_1, x_2))$.

We now describe the properties of β_1: First, $\beta_1(L_1, x_2) = L_1$ because $S(x_2, L_1) = L_1$. Second, the fact that $S(x_2, w_1)$ is an integrated cdf implies $\beta_1(x_1, L_2) = x_1$ and $x_1 > \beta_1(x_1, x_2) > L_1$ away from (L_1, L_2). Third, suppose for simplicity that S is partially differentiable in w_1, and implicitly differentiate of the first-order condition in equation 2: $\frac{\partial \beta_1}{\partial x_2} = -\frac{\partial S}{\partial x_2} \Big/ \left(1 + \frac{\partial S}{\partial w_1}\,|_{w_1=\beta_1}\right) < 0$ away from (L_1, L_2) because the numerator is negative by convexity of S in x_2 (and the denominator is positive by first-order regularity of S). Similarly, $\frac{\partial \beta_1}{\partial x_1} = \left(1 + \frac{\partial S}{\partial w_1}\,|_{w_1=\beta_1}\right)^{-1} > 0$, so $\beta_1(x_1, x_2)$ is increasing in x_1. This confirms that the equilibrium is strictly monotone. The curvature of $\beta_1(x_1, x_2)$ is ambiguous. *QED Step 3. QED Theorem* 1

References

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

Bergemann, Dirk and Johannes Horner. 2010. Should Auctions be Transparent? Cowles Foundation Working Paper.

Black, Jane, and David de Meza. 1992. Systematic Price Differences between Successive Auctions are no Anomaly. *Journal of Economics and Management Strategy* 1 (4): 607–8.

Board, Simon 2009. Revealing information in auctions: the allocation effect. *Economic Theory* 38:125–135

Budish, Eric. 2008. Sequencing and Information Revelation in Auctions for Imperfect Substitutes: Understanding eBay's Market Design. Working paper. Harvard University.

Bulow, Jeremy and Paul Klemperer. 1996. Auctions versus Negotiations. *American Economic Review.* 86(1), 180-194.

Che, Yeon-Koo, and Ian Gale. 1998. Standard Auctions with Financially Constrained Bidders. *Review of Economic Studies* 65 (1): 1–21.

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

Engelbrecht-Wiggans, Richard. 1994. Sequential Auctions of Stochastically Equivalent Objects. *Economic Letters* 44: 87–90.

Gale, Ian L., and Donald B. Hausch. 1994. Bottom-Fishing and Declining Prices in Sequential Auctions. *Games and Economic Behavior* 7: 318–31.

Haruvy, E. and P. T. L. Popkowski Leszczyc. 2010. Search and Choice in Online Consumer Auctions. *Marketing Science* 29(6): 1152-1164.

Holt, Charles A. and Roger Sherman. 1994. The Loser's Curse. *American Economic Review* 84(3): 642-52.

Jehiel, Philippe and Benny Moldovanu. 2000. Auctions with downstream interaction among buyers, *RAND Journal of Economics* 31 : 768–791.

Jofre-Bonet, Mireia, and Martin Pesendorfer. 2003. Estimation of a Dynamic Auction Game. *Econometrica* 71: 1443–89.

Kambil, Ajit and Eric van Heck, 1998. Reengineering the Dutch flower auctions: a framework for analyzing exchange organizations. *Information Systems Research* 9 (1): 1–19.

Katzman, Brett. 1999. A Two-Stage Sequential Auction with Multi-Unit Demands. *Journal of Economic Theory* 86: 77–99.

Kim, Kyoo il and Joonsuk Lee. 2009. A Structural Analysis of Wholesale Used-Car Auctions, working paper. University of Minnesota.

Krishna, Vijay. 2002. *Auction Theory.* Academic Press of Elsevier Science, London, UK.

Levin, Jonathan. 2011. The Economics of Internet Markets. Forthcoming, *Advances in Economics* *and Econometrics*, Edited by D. Acemoglu, M. Arellano, and E. Dekel.

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

Milgrom, Paul R., and Robert Weber. 2000. A Theory of Auctions and Competitive Bidding II. Working paper in 1982, now published in: Klemperer, P. (ed.), *The Economic Theory of Auctions*. Cheltenham, UK: Paul, Edward Elgar Publishing.

Pesendorfer, Wolfgang and Jeroen M. Swinkels. 1997. The Loser's Curse and Information Aggregation in Common Value Auctions. *Econometrica* 65(6) : 1247-1281

Reiß, J. Philipp and Schöndube, Jens Robert. 2008. First-price equilibrium and revenue equivalence in a sequential procurement auction model. *Economic Theory*, forthcoming.

Roth, Alvin E. and Axel Ockenfels. 2001. Last-Minute Bidding and the Rules for Ending Second-Price Auctions. *American Economic Review.* 92(4), 1093-1103.

Tadelis, Steven and Zettlemeyer, Florian. 2015. Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment. American Economic Review 105(2), 886-905.

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

Zeithammer, Robert. 2005. Sequential Auctions with Information about Future Goods. Working Paper. University of Chicago.

Zeithammer, Robert. 2006. Forward-Looking Bidding in Online Auctions. *Journal of Marketing Research* 43(3): 462–76.