Internet Auctions with Many Traders

Michael Peters and Sergei Severinov
University of Toronto and University of Wisconsin

First Version Feb/01 This version April 4, 2002

Abstract

A multi-unit auction environment similar to Ebay is studied. Sellers who wish to sell a single unit of a homogenous good set reserve prices for their own independently run auctions. Buyers who hope to acquire a single unit bid as often as they like in a dynamic second price auction. When the number of buyers and sellers is large but finite, there is a Bayesian equilibrium for this completely decentalized trading procedure in which the ex post efficient set of trades occurs at a uniform trading price. Remarkably, the strategy rules that buyers and sellers use in this equilibrium are very simple. They do not depend in any way on beliefs, or on the number of buyers and sellers.

Internet auctions and business to business trading sites play have captured a significant proportion of trade in modern markets. For example, eBay had 18.9 million registered users as of April 2001. Total sales volume at eBay in 2000 was $ 5 bln. and this number is expected at least to double in 2001 (see Morneau (2001)).The volume of trade at business-to-business marketplaces has grown from $5 bln. in 1999, to $43 bln. in 2000 and, according to ActivMedia Research, is expected to reach $263 bln. in 2001 (see Wilson and Mullen (2000), Morneau (2001)). Though computerized bidding for spectrum auctions, for example, has been heavily influenced by economists and economic theory, the more popular trading institutions like eBay have emerged on their own.

As is often the case, economic theory and practise diverge considerably concerning the way that these large exchanges should be organized. Many of the goods traded on internet auction markets are commodities, like computer processors ((Zheng 2001) or (Carare 2001)) or software. An obvious suggestion is that these should be traded on a double auction. Double auctions are *incentive efficient* Wilson (1985) and become *ex-post efficient* quickly as the number of buyers and sellers gets large (Rustichini, Satterthwaite, and Williams 1994).

When goods are less perfectly substitutable, like coins ((Bajari and Hortacsu January, 2000)) or antiques ((Roth and Ockenfels 2000)), goods should be sold using a Vickrey mechanism. In the case where buyers want a single unit and sellers all have one unit to sell, the Vickrey auction asks buyers to describe their willingness to pay for *each* of the items on offer. Sellers give their ask prices. The auctioneer then chooses an allocation of goods to buyers that maximizes total apparent surplus (given the asks of sellers) and sets prices so that each buyer's private surplus is equal to the total surplus when he is included and the total surplus when his good is reallocated somewhere else. A seller's bid double auction (described more formally below) is a Vickrey mechanism for the case of identical goods, so at least in the case where all goods are identical, the Vickrey mechanism is incentive efficient. Whether this is true in the more general case with different goods is not known, but the Vickrey mechanism does maximize total seller revenues within the class of all auctions that are efficient in the sense that they maximize surplus given sellers announced costs (Krishna and Perry 1998).

Vickrey mechanisms have desirable incentive properties in the sense that buyers do at least as well describing their true preferences to the Vickrey auctioneer as they do by lying no matter what other buyers do.[1] Secondly, there are algorithmic procedures that can be used to determine the Vickrey prices and allocations. One such procedure is described for the case where each bidder wants only a single good in (Roth and Sotomayor 1990). An algorithmic procedure that leads to the Vickrey allocation in more general situations is given by (Ausubel and Milgrom 2001).

Though there are business to business exchanges that use centralized trading mechanisms like the Vickrey mechanism[2], the successful internet trading sites are organized more along the lines of a competing auction market (Peters and Severinov 1997, McAfee 1993, Peters 1997). On eBay, for example, sellers conduct their own auctions in which they chooses minimum bids and possibly secret reserve prices. The auctions are otherwise ascending second price auctions. Buyers submit bids to a particular seller's auction by giving a proxy bid to a robot who then sets the standing bid in the auction equal to the second highest bid it has received.[3]

The critical difference between this kind of market design and the Vickrey procedure is that buyers are given the responsibility of allocating themselves, and adjusting prices during the procedure. This gives them much greater scope to manipulate the pricing algorithm than they have when they simply submit preferences to a Vickrey auctioneer. They could misrepresent their preferences,

---

[1] There is no reason why sellers should reveal their costs truthfully. A natural question to ask is whether large numbers will suffice to induce sellers to set reserve prices equal to their true costs as is the case in double auctions. We are not aware of any results of this kind.

[2] For example, the NTE truck exchange (http://www.nte.net) allows truckers with excess carrying capacity to match with shippers who have small and unusual packages to send. The exchange does not use auctions, but instead uses what they call 'market based pricing' to complete the trade. Apparently the 'commodities' used in this exchange are so narrowly defined that mechanisms similar to a double auction that use only demand and supply for a single 'commodity' to determine the price attract too few buyers and sellers, so supply and demand information for related routes and times is aggregated and used to generate market prices.

[3] It is actually a bit difficult to cull details of the price setting process out of the information given on eBay's website. An approximate description of the price setting procedure, along with a program that lets you watch the bidding, is available at http://manitas.economics.utoronto.ca/~peters/ebay_intro.php. A good description of the procedure is given in (Roth and Ockenfels 2000).

as in the Vickrey mechanism. They can do this in a way that is dependent on the history of the bidding process (which they can't in Vickrey). They can even manipulate the price adjustment procedure itself by bidding in ways that are inconsistent with algorithmic price adjustment rules. It is straightforward to show (we do so below) that 'reporting preferences truthfully' then allowing the proxy bidder to adjust prices as if the mechanism were a Vickrey procedure may be strictly dominated by a more sophisticated strategy.

The purpose of this paper is to show that despite this, there is a decentralized bidding process, similar (but different in some important ways) to what happens on eBay, for which there exists a full sequential Bayesian equilibrium which supports efficient trade at Vickrey prices in one special, but nonetheless interesting, environment. The bidding procedure adapts the process used in the literature on competing auctions (for example (Peters and Severinov 1997)) to the case where bidding is costless. It mimics eBay by having sellers to run independent ascending auctions where they are free to set any reserve price that they wish. Buyers bid sequentially, choosing both the location and amount to bid.

To emphasize the distinction between our work and the previous literature, the price adjustment procedures described by, for example, (Ausubel and Milgrom 2001) or (Roth and Sotomayor 1990) are *imposed* on the market. Our price adjustment procedure is *derived* by finding a perfect Bayesian equilibrium corresponding to the bidding rules. We do not need to invoke the incentive properties associated with Vickrey allocations, because the bidding strategy we describe is fully sequentially rational. The remarkable part of our theorem is that despite the fact that the actual price adjustment is fully endogenous, it *looks* very much like an algorithmic price adjustment procedure in the sense that the sequentially rational bid at each information set is independent of bidders beliefs, and depends only on the current array prices and possibly the outcome of the last bid with each seller. Bidders could easily write their own proxy bidding routines to implement the strategy for them. Alternatively, eBay could simply provide the proxy bidding routine to the bidders without worrying that the bidders would try to outwit it or interject their own routines.

A second remarkable property of our theorem, is that the sequential Bayesian equilibrium for the bidding process among buyers puts sellers in exactly the same strategic situation when they set their reserve price, as they would encounter in a seller's bid double auction. From (Rustichini, Satterthwaite, and Williams 1994), this should imply that sellers will want to set their reserve prices close to their true costs when there are many traders with independent valuations. We cannot impose their theorem directly for a couple of reasons, so we provide a proof of the fact that when the number of buyers and sellers is large enough (but still finite) there exists a full equilibrium for our bidding procedure in which sellers set their reserve prices equal to their true costs. As a consequence, the bidding process we describe yields a fully ex post efficient equilibrium when there are many traders.

1 The Model

There are n sellers and m buyers trading in a market. The number of buyers and sellers is arbitrary in our model, and we do not assume that there are more buyers than sellers. Each seller has one unit of a homogeneous good, while each buyer has an inelastic demand for one unit of this good. Buyers' valuations and sellers' costs are private information and are distributed on the grid $\mathcal{D} \equiv \left\{\underline{d}, \underline{d}+d, \underline{d}+2d, \ldots \overline{d}\right\}$ that has a step size $d > 0$. Let $F(\cdot)$ and $G(\cdot)$ be the probability distributions from which buyers' and sellers' valuations respectively are drawn. Our results on equilibrium bidding behavior by buyers are independent of whether or not buyers' valuations are correlated, though our equilibrium for the sellers' part of the game relies on independence. A buyer with valuation v who wins a single auction at a price p gets surplus $v-p$. In this sense we consider a purely *private value* model. A buyer who wins more than one auction gets no additional utility from the additional units of output (so his payoff will fall because he has to pay for the additional units). A sellers with cost c who sells at price p get surplus $p-c$.

Trade is organized in the following way. At first, sellers simultaneously announce reserve prices for their auctions. Thereafter buyers arrive sequentially. When a new buyer arrives, he submits a bid at whichever of the sellers' auctions he likes best. Buyers are required to submit bids in the grid $\mathcal{D}$.[4] The grid size d could be thought of as the minimum bid increment. When the seller receives a bid, she publishes a number called her *standing bid* which is equal to the second highest bid that she has received, or her reserve price if she has not received more than 1 bid. Each seller immediately updates her standing bid announcement when her standing bid changes. As in most existing on-line auctions, we also assume that the identity of the winning bidder is known at all times[5].

We assume throughout that the second-highest bid means second highest bid submitted by a distinct bidder. The standing bid is assumed to remain unchanged if the high bidder revises his bid. A new bid submitted at a seller's auction must always exceed that seller's current standing bid. If two or more bidders have submitted the same high bid, then the buyer who was the first to submit this bid is declared the high bidder. The standing bid in this case is equal to the high bidder's bid.

After a new bid is submitted, each buyer in order of his or her entry into the market is given the opportunity either to submit a new bid (not necessarily with the same seller) or pass. Once each buyer in the market chooses to pass, a new buyer enters. After all buyers have entered the market, the bidding process continues as bidders update their bids one after another. The order of bidding at this stage is the same as the order of entry with the last bidder followed by the first bidder and so on. The bidding continues until all buyers pass. Then

---

[4] This assumption is natural in view of our interpretation that the grid on which the traders' valuations are distributed is determined by the minimal monetary unit.

[5] As will be show later, this assumption, or more precisely, the observability of the change in the identity of the winning bidder is important for the uniform price result. When such changes are not observable, price dispersion may occur

the high bidder at each seller trades at the final standing bid with that seller.

These rules approximate the trading rules on Ebay and on the Amazon auction site (which differs from Ebay primarily in that auctions do not have a definite ending date (see Roth and Ockenfels (2000))). The key property of the auction rules is that they generate a type of the second price auction in which the high bid is never observed.

Despite the second price nature of the auction mechanism, the presence of multiple auctions implies that it is not a dominant strategy for buyers to bid their true valuations when they start bidding (despite the advice offered on the Ebay website). To see this, consider the data in Figure 1. For the moment, ignore the point b_3, and suppose that there are two buyers with true valuations b_1 and b_2 and two sellers who announce reserve prices s_1 and s_2. Assume that the buyer with valuation b_1 enters first and expects the other buyer to bid his true valuation wherever he decides to bid.

To show that it is not a dominant strategy, we only need to show that valuation bidding can be strictly improved upon for *some* strategy of buyer 2, so suppose that buyer 2 is expected to bid his valuation where the standing bid is lowest. Buyer 1 could bid her valuation with seller 2. She should then expect buyer 2 to bid his valuation with seller 1 who has a lower reserve price. No matter what bid buyer 2 submits, buyer 1 will trade at price s_2. Buyer 1 can strictly increase her expected payoff by initially bidding s_2 at seller 1, provided she believes that buyer 2 has a valuation below s_2 with strictly positive probability.

So if buyer 1 is going to start off bidding her valuation, she will bid with seller 1. When buyer 2 enters, the standing bid with seller 1 remains at s_1 because seller 1 has yet to receive a second bid. So by the rule that buyer 2 is supposed to be using, buyer 1 should expect buyer 2 to bid against her provided he bids at all.

If buyer 2 bids at or below s_2, then buyer 1's initial strategy works out fine, and she trades at a price at or below s_2. If buyer 2 submits any bid above b_2, then buyer 1 will be displaced as high bidder, and will trade at price s_2 with seller 2 - again a reasonable outcome. The difficulty arises when buyer 2 submits a bid strictly above s_2 but strictly below b_2. In this case, buyer 1 will pay a price strictly above s_2. It is not hard to see that if buyer 1 starts with a bid s_2 with seller 1 (instead of b_1), then everything works out the same way, except that when buyer 2 bids something between s_2 and b_1, buyer 1 is displaced as high bidder and trades at price s_2 with seller 2. Deviating to the price s_2 ensures that buyer 1 will never have to pay a price above s_2. This strictly improves upon the bid b_2 provided buyer 1 believes that buyer 2 has a valuation between s_2 and b_1 with strictly positive probability.

2 Efficient Bidding

The advantage of sequential bidding is that buyers whose valuations are high, but who had the bad luck of bidding against another buyer with an even higher

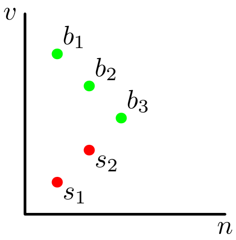

Figure 1:

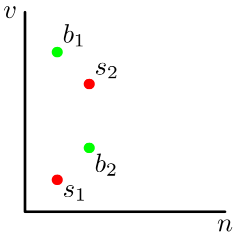

Figure 2: Inefficient Equilibrium

bid, have an option to bid again elsewhere. Unfortunately, this option is not sufficient to guarantee that, conditional on seller's reserve prices, the efficient trades are carried out. The following example demonstrates this.

In Figure 2, given sellers' reserve prices s_1 and s_2 the efficient outcome is for buyer 1 to trade with seller 1. However, consider the following strategies. If a buyer finds that no other bids have been submitted and his valuation is at least s_1, then he submits a bid equal to s_2 with seller 1. If there is a bid at seller 1, then the buyer bids his valuation with seller 2 provided his valuation is at least as large as s_2, and refrains from bidding otherwise. This strategy is optimal for the buyer who enters first because she ends up trading with seller 1 at price s_1. Following this strategy is also optimal for the buyer who enter last, because he believes (correctly) that the buyer who has already entered has bid s_2 at seller 1. Therefore, there is an equilibrium in which both buyers use this strategy. Given the data in Figure 2, if buyer 2 enters first then in this equilibrium he will trade with seller 1, and then buyer 1 with a higher valuation will trade with seller 2. Thus, equilibrium outcome of the bidders' game could be inefficient.

It is easy to demonstrate that in this example, as well as in many other ones, the dynamic bidding game has multiple equilibria. We will not attempt to characterize all of them. Rather, our objective is to try to identify bidding equilibria that have nice properties, especially, the ones that are efficient conditional on the announced reserve prices.

To begin, let $\mathbf{b} = \{b_1, \ldots b_m\}$ be the vector of buyers' valuations, and let $\mathbf{s} = \{s_1, \ldots s_n\}$ be the vector of sellers' reserve prices. Without loss of generality, assume that sellers are indexed in such a way that $s_1 \leq s_2 \cdots \leq s_n$, while buyers are indexed the opposite way so that $b_1 \geq b_2 \cdots \geq b_3$. Following Satterthwaite and Williams (1089), let $\mathbf{v} = \{v_1, v_2, \ldots v_{m+n}\}$ be vector with entry v_i equal to i-th lowest value among buyers' valuations and sellers' reserve prices. If sellers' reserve prices were equal to their true valuations, then the efficient set of trades occurs if buyers whose valuations exceed v_m trade with sellers whose reserve prices are equal or less than v_m. To see why, note that initially there are n sellers holding one unit of output and m buyers. Therefore, after all trades are completed exactly n traders will own the good and the other m will not. Efficiency requires that the traders who end up without the good have the lowest valuations or costs. The price v_m is the lowest price that clears the market given the demand and supply schedules generated by $\mathbf{b}$ and $\mathbf{s}$. It is the price that would be chosen by the auctioneer is a 'seller's bid double auction' ((Rustichini, Satterthwaite, and Williams 1994)).

There is yet another way to think about the price v_m. Let V be the maximal *apparent* gains to trade between buyers and sellers - i.e., the sum of the differences between the valuations of the buyers who trade and the reserve prices of the sellers who they trade with in an efficient matching of buyers to sellers. Let V_{-i} be the corresponding maximal apparent gains to trade when buyer i is left out of the allocation. The Vickrey price for buyer i is the price p_i such that $b_i - p_i = V - V_{-i}$. It is not hard to see that if buyer i trades in an efficient allocation, then absent buyer i the best thing to do with the unit of output is to give it to the buyer or seller who has valuation v_m. So $V - V_{-i}$ must equal the difference between buyer i's valuation and v_m. In other words, the Vickrey price for every trade is equal to v_m.

We now define the symmetric strategy σ^* by specifying how each buyer should bid when his (her) turn comes. To help describe the strategy, make the natural definition and say that a bid is *successful* if the bidder who makes it becomes high bidder, and that the bid is *unsuccessful* otherwise.

Definition 1 *The symmetric strategy σ^* is defined as follows: if it is the buyer's turn to bid then*

(a) *if the buyer is the current high bidder at any auction, or if the buyer's valuation is less than or equal to the lowest standing bid, the buyer should pass;*

(b) *otherwise, if there is a unique lowest standing bid, the buyer should submit a bid with the seller offering this lowest standing bid. The bid should be equal to the lowest value on the grid that exceeds this low standing bid;*

(c) *otherwise, if more than one seller has the lowest standing bid, the buyer should submit the same bid as in (b) with equal probability at each such seller where either the seller has not received a bid, or the last bid the seller received was unsuccessful. If the last bid was successful with all sellers holding the lowest standing bid, then the buyer should bid with each of them with equal probability.*

Before proceeding, it might help to visualize the path generated when buyers use strategy σ^* in a simple demand-supply style diagram. We refer again to Figure 1, where the valuations and reserve prices of three buyers and two sellers are shown. This is simply an example designed to illustrate the way the strategies work, so for convenience we can assume that the grid of feasible bids coincides with the set of sellers' reserve prices s_1, s_2 and buyers' valuations b_3, b_2, b_1. Assume further than buyers enter in reverse order of their valuations, so buyer 3 with valuation b_3 enters first, then buyer 2 then buyer 1. According to strategy σ^*, when buyer 3 enters he bids s_2 with seller 1 because seller 1 initially has the lowest standing bid and s_2 is the lowest valuation on the grid that exceeds s_1. This bid will be successful but it will have no effect on seller 1's standing bid. Buyer 2 will bid the same amount with seller 1, since seller 1's standing bid still has the lowest standing bid.

The bid by Buyer 2 is unsuccessful, but it makes the standing bid with seller 1 rise to s_2. Observe that unsuccessful bids will always change a seller's standing bid, since buyers must submit bids above the current standing bid. As buyer 3 was the first to submit this bid, he remains the high bidder and passes by (a), so buyer 2 has the chance to submit a new bid. Now both sellers have the same standing bid. The last bid with seller 1 was not successful, and seller 2 has not yet received a bid. The lowest value that exceeds the standing bid s_2 is b_3, so by (c) buyer 2 will bid b_3 with equal probability at seller 1 (case (a)) or seller 2 (case (b)). In either case, buyer 2's bid will be successful and it will not affect either of the standing bids. Note that a successful bid may or may not change a seller's standing bid.

Consider case (a) first. Since buyer 2 displaces buyer 3, buyer 3 immediately submits a bid equal to b_3 with seller 2, since seller 2 has yet to receive a bid in this case. The bid is successful, and again, this bid does not affect either of the standing bids which remain equal to s_2. When buyer 1 enters he finds that the last bid with both sellers was successful, so he will randomize. Whichever sequence he chooses, the bid b_3 will be unsuccessful and will raise the standing bids at both sellers will rise to b_3. Then buyer 1 will bid b_2 choosing randomly between the two sellers (since the last bid was unsuccessful at both locations). If buyer 1 bids first with seller 1, then he will displace buyer 2, who will, in turn, bid b_2 at seller 2 and become a winning bidder there. Bidding will then stop, and buyers 1 and 2 will trade at price b_3. If he bids first with seller 2, his bid will be successful and all bidding will stop.

Now consider case (b). When buyer 1 enters, he will submit bid b_3 at seller 1 and displace bidder 3 without raising the reserve price above s_2. Buyer 3 will then submit bid b_3 at sellers 1 and 2 in random order, which will raise the standing bid at both sellers to b_3, but bidder 3 will still not be a winning bidder. The bidding will then stop and buyer 1 and 2 will trade.

This example conveys the essential idea. Buyers bid up prices with each seller as slowly as possible. For this reason, high valuation buyers are never trapped into paying higher prices if another high-valuation buyer accidentally bid against them. In this example, the efficient trades occur and both sellers trade at a uniform price equal to buyer 3's valuation. The randomness on the equilibrium path makes possible different profiles of winning bids and different pairwise matching between buyers and sellers who trade. However, the uniform trading price is uniquely determined by the profiles of buyers' valuations and sellers' reserve prices. These properties of the strategy rule σ^* are quite general, as the following theorem demonstrates.

Theorem 2 *The outcome in which all buyers use the strategy* σ^* *is a perfect Bayesian equilibrium. For each array of valuations and reserve prices* $\mathbf{v}$*, each buyer whose valuation is above* v_m *trades with some seller whose reserve price is no larger than* v_m*. All trades occur at the price* v_m*.*

The proof of this and the other main theorems in the paper is included in the appendix.

A couple of comments are in order at this point. The environment that we consider here is a special case of the one to one matching environment discussed in (Roth and Sotomayor 1990). Their environment differs from ours in that buyers may have different values for the goods offered by different sellers and we do not allow this. Suppose that buyers follow σ^* from the beginning of the bidding process. At any stage in the process, if there are strictly more buyers willing to pay the lowest standing bid than there are sellers offering this price, then the set of goods offered at this lowest standing bid constitutes what Roth and Sotomayor call an *overdemanded set.* The subsequent bidding will raise each of the standing bids of sellers in that set by one grid point. So in this sense the bidding rule implements the algorithm discussed by Roth and Sotomayor. The content of Theorem 2 is to show that there is at least one perfect Bayesian equilibrium for which the algorithm is not subject to strategic manipulation at any time during the course of the procedure. What makes the argument difficult is the fact that the bidding rule has to deal with situations that could never occur under the simple application of the Roth Sotomayor procedure.

One could also apply the (Ausubel and Milgrom 2001) ascending proxy auction idea to our environment. They study a first price procedure which takes place in a series of rounds. At the end of each round, the high bids with each seller are known. In our simple environment, bidders would are not high bidders at any auction would simply choose one of the sellers where price is lowest and bid the minimal price above the current high bid. The auction ends when no new bids are submitted. We are interested in a much different kind of incentive issue than are (Ausubel and Milgrom 2001), which is why we focus on the simpler economic environment. It is likely the arguments we describe could be applied to the first price procedure as well. We chose to study the second price procedure for practical reasons - it more closely resembles the procedure that is employed on eBay.

According to strategy σ^*, all active buyers focus their bids on seller(s) with the lowest standing bids irrespective of buyers' valuations, beliefs and the standing bids. The active bidders continue to bid up the standing bid with this group of sellers (which could consist of only one seller) until their standing bids reach the level of the standing bids or reserve prices at the next group of sellers, i.e. the second lowest level. Then bidding continues with sellers from both groups until the lowest standing bid reaches the level of reserve price or standing bid at the next group of sellers, and so on. This process continues until all remaining active bidders become high bidders with different sellers.

It is notable that strategy σ^* contains the same set of rules on the equilibrium path as well as off the equilibrium path. But, although σ^* generates a relatively simple path, the proof that playing σ^* constitutes an equilibrium is not straightforward because the dynamic bidding game is sufficiently complex.

We provide the proof by showing that a buyer who at any stage of the bidding game submits a higher bid than what is prescribed by σ^*, or who bids above his valuation can at best raise the trading price. At the same time, a buyer who bids as if his valuation is lower than what it is in reality, can sometimes lower the trading price, but only by giving up a desirable trading opportunity.

This line of argument is similar in spirit to the one that is used to show that bidding the true valuation is an equilibrium in a standard static second price auction. But the analogy is only approximate.[6] A substantially more complex argument is needed because of the dynamic nature of the bidding process and presence of multiple auctions. To compute the payoff that a buyer gets by following σ^* from the outset of the game, as well as the payoffs that he could get by deviating, we need to consider arbitrary information sets and characterize the outcomes that occur after the play of the game has reached them. In particular, we need to show that playing σ^* remains optimal for a buyer no matter what information set is reached, if in the continuation all other buyers follow σ^*.

An important property of the equilibrium where buyers use σ^* is absence of price dispersion in the final outcome: all trades are executed at the same price. Rule (c) in σ^* plays the crucial role in achieving this. This rule describes buyers' behavior when there are several sellers with the lowest standing bid, and so the buyer faces the uncertainty regarding the winning bids at these sellers.

The essence of rule (c) is that it allows the buyer to identify which of the sellers with the same standing bid have lower winning bids. Bidding only at such sellers is optimal for the buyer, because then with some probability the buyer will trade at a lower price. At the same time, such bidding ensures that standing bids rise uniformly and eliminates the possibility of price dispersion.

Rule (c) tells the buyer to focus on such sellers with lowest standing bids where there had been less action at the current standing bid, or more precisely, where the last submitted bid was unsuccessful in the sense that the bidder who submitted it did not succeed in displacing the high bidder at the auction. On the path generate by our strategies, a successful bid will be strictly above the standing bid and will displace the lead bidder without changing the standing bid. An unsuccessful bid will raise the standing bid without displacing the winner. Thus, when two standing bids are the same, the one where the last bid was successful will have a larger high bid.

---

[6] As shown above, bidding one's valuation is not an optimal strategy in the dynamic bidding game.

For each array **v** of buyers' valuations and sellers' reserve prices, theorem 2 uniquely specifies the price at which all trades will be completed. Precisely, the trading price will be equal to v_m which is either the highest *valuation* among buyers whose bids are unsuccessful, or the highest reservation price among sellers who trade, depending on the actual array of valuations and reserve prices. At the same time, the randomization on the equilibrium path (buyers randomize between the sellers among whom they are indifferent) implies that certain aspects of the outcome will also be random. First, the final matching between sellers and buyers who trade is going to be random. Second, if there are $k > 1$ traders whose valuations or costs are equal to v_m, then any number of them between zero to $k-1$ may trade in equilibrium.

To better understand whether buyers and sellers with valuations and costs equal to v_m trade, let us divide the set of buyers into three groups. For any array **v**, let M_1 be the set of buyers whose valuations are *strictly lower than* v_m, M_2 be the set of buyers who have valuations *exactly equal to* v_m, while M_3 be the set of buyers whose valuations are *strictly higher than* v_m. Similarly, let N_1, N_2, and N_3 be the sets of sellers who set their reserve prices below, equal to and above v_m respectively. Let m_i (n_i) be the number of buyers (sellers) in the set M_i (N_i). Theorem 2 says that buyers in M_3 will surely trade, and that buyers in M_1 and sellers from N_3 will not trade.

Corollary 3 *Let v_m be the m-th lowest element in the array* **v** *of buyers' valuations and sellers' reserve prices. If all buyers use the strategy σ^* in the bidding game, a seller who sets reserve price s s.t. $s < v_m$ trades for sure. The number of sellers with reserve price equal to v_m who trade is between*

$$\max\left[0, m_3 - n_1\right] \quad \textit{and} \qquad \min\left[n_2, m_3 - \min\{0, n_1 - m_2\}\right]$$

The number of buyers with valuation equal to v_m who trade is between:

$$\max\left[0, n_1 - m_3\right] \quad \textit{and} \qquad \min\left[m_2, n_1 - \min\{0, m_3 - n_2\}\right]$$

Proof: see the appendix.

Since all sellers from N_2 (buyers from M_2) are identical, they have the same chances of trading. Therefore the probability that a seller from N_2 (buyer from M_2 trades lies in the interval with boundaries that are derived by dividing the corresponding boundaries on the number of sellers from N_2 (buyers from M_2) who trade over N_2 (M_2).[7]

---

[7] By modifying the definition of strategy σ^* appropriately, we can support equilibria in which the number of sellers from N_2 who trade is equal to the lower or upper bound established in corollary 3. To obtain the lower bound, modify σ^* in the following way. When a buyer faces the choice between several sellers who have the lowest standing bid and who are equivalent with respect to rule (c), he bids at first with those sellers whose original reserve prices were lower than their current standing bid. Then the buyers in M_3 will submit bids with sellers in N_1 before they will consider the sellers in N_2. When the opposite modification on the strategy σ^* is imposed, we can support an equilibrium that reaches the upper bound on the number of sellers from N_2 who trade.

The central implications of the corollary 3 is that the outcome of the bidding game is efficient: traders who are left without the good at the end of the day are the ones who have the lower valuations and costs. Note that no matter how many sellers from N_2 and buyers from M_2 trade, this has no effect on the other traders and the efficiency of the outcome.

We view the presence of randomness in the final outcome as a strength of our model. In large decentralized markets it is unrealistic to expect all aspects of the bidding process to be entirely deterministic. Yet, we demonstrate that despite the presence of randomness on the equilibrium path, the final price will be uniform and independent of the actual path of bidding.

It is also worth pointing out the connection between our result and equilibrium in the double auction market. Consider a double auction with uniform trading price set equal to m-th lowest value among buyers' bids and sellers' asks. Satterthwaite and Williams (1089) refer to it as 'seller's bid double auction'. They show that buyers would bid their true valuations in it and therefore, conditional on the seller's ask prices (reserves), the trading price will be equal to v_m. Bidding his valuations is optimal for a buyer because his bid can only affect the trading price if the buyer fails to trade. The usual second price logic then implies then applies. Thus, the dynamic bidding game studied here has an equilibrium that generates an outcome identical to the outcome in a seller's bid double auction when sellers' reserves prices are the same set in both markets.

Note that neither the description of strategy σ^* nor the proof that σ^* is a best reply depends in any way on the distribution of valuations, or the number of buyers and sellers in an auction.[8] For example, if there is only a single seller, the trade will occur at the reserve price if there is only a single buyer whose valuation is above the reserve, and at the second highest valuation otherwise. This is the same as the outcome in a second price auction. Playing strategy σ^* remains an equilibrium when there is a large number of sellers. On the other hand, as in the seller's bid double auction, a seller's behavior in our mechanism will typically depend on her beliefs and the number of traders. We turn to the this issue in the next section.

To prove the strong incentive properties of the bidding procedure, we have focussed on a particularly simple economic environment. In particular, we have focussed on the case where all goods are identical. This limits the applicability of our result. On eBay, for example, it is possible to restrict attention to commodities that would appear to be very close substitutes. For example, (Zheng 2001) focuses on the market for 800 MHz computer chips. Even within this narrowly confined market, there are important differences between the objects offered for sale by each seller. One difference in particular is that auctions for these goods do not end at the same time, as we have assumed in this paper. It is not hard to show that the bidding rule σ^* does not always lead to Vickrey prices when goods

---

[8] This statement has to be qualified slightly since we impose restrictions on beliefs off the equilibrium path, for example to ensure that no buyer believes that the high bidder in a any auction has a valuation below the standing bid in that auction. This restriction may seem unreasonable if valuations are correlated and prior beliefs conditional on the buyer's own valuation are inconsistent with this.

are not identical, assuming that all bidders follow it. This raises a number of difficult problems. The most obvious is simply whether there are decentralized mechanisms (in the sense that all price adjustment and allocation decisions are made by buyers and sellers themselves) which have desirable perfect Bayesian equilibria in the case where goods are not identical (or where agents want to buy many goods). The second is to characterize the perfect Bayesian equilibria for the existing mechanism, again when goods are not identical.

3 Sellers' Strategies and the Efficiency Result

The results of the previous section indicate that buyers' equilibrium strategies guarantee that the efficient set of trades occurs when sellers set reserve prices equal to their true valuations. Buyers behavior is efficient because buyers affect the trading price only when they forego attractive trading opportunities.

At the same time, the outcome will not necessarily be efficient if sellers set reserve prices that are different from their true costs. So we turn to an examination of sellers' behavior in this section.

It may not be optimal for a seller to set a reserve price equal to her true valuation, because her reserve price may affect the trading price in some cases. For example, consider the situation depicted in Figure 1 without buyer 3 (the one with the lowest valuation). From the Figure it is clear that, if buyers use strategy σ^*, the uniform trading price will be equal to the higher reserve price as long as it is below b_2. Sellers, of course, do not know b_2. Thus, seller 2 could raise the trading price with a strictly positive probability by raising his reservation price above s_2. Clearly, the cost of increasing her reserve price to seller 2 is that she would fail to trade if either buyer has valuation between s_2 and her new higher reserve price.

This trade-off is similar to the one which traders face in a standard double auction. The rate of convergence of the optimal reserve prices to their true costs in a double auction has been analyzed by Satterthwaite and Williams (1089) in the case where the costs and valuations are independently and continuously distributed over an interval. In the analysis below, we will also maintain the independence assumptions. However we consider a finite grid of valuations which allows us to get a somewhat tighter result.

Thus, assume that buyers' valuations and sellers' costs are distributed independently. Let $f(p) \equiv F(p) - F(p-1)$ and $g(p) \equiv G(p) - G(p-1)$ denote the probability that a buyer's valuation and a seller's cost respectively are equal to p exactly. Let $\bar{g} = \min_{p \in \mathcal{D}} g(p)$ and $\bar{f} = \min_{p \in \mathcal{D}} f(p)$. We further assume that $\bar{g} > 0$ and $\bar{f} > 0$.

We consider a sequence of markets that get larger as the number of traders increases. For simplicity, we hold the ratio of the number of buyers to the number of sellers constant at $k > 0$ i.e., $m = kn$ where m is the number of buyers and n is the number of sellers.

The main result of this section is the following theorem which establishes that setting a reserve price equal to the true cost constitutes an equilibrium strategy for sellers when the number of traders in the market is sufficiently large.

Theorem 4 *Suppose that every seller except seller z sets her reserve price equal to her true cost and buyers follow the strategy σ^*. If m and n are sufficiently large, it is optimal for seller z to set reserve price equal to her true cost.*

The proof of this theorem is provided in the appendix. It demonstrates that, provided that the number of traders is sufficiently large, a seller with cost c obtains a higher expected payoff by setting reserve price equal to $p-d$ rather than p for $p > c$. The two expected payoffs are the same if the trading price is either above p or below $p-d$. Thus, we need to focus on the situations where seller j's choice between p and $p-d$ affects the trading price, and where the trading price is equal to $p-d$ or p irrespective of seller j's choice. The latter situation occurs if several other sellers post the reserve price equal to the trading price. The probability of this event is zero when the costs are distributed continuously. Yet, with a discrete set of valuations it occurs with a positive probability and has an important effect. When several sellers are tied at the trading price, a seller may fail to trade even if her reserve price is equal to the trading price, because of the competition from the other sellers posting this reserve price.

Precisely, seller j gets different expected payoffs by setting reserve price equal to $p-d$ or p in the following four cases:

1. the trading price is equal to $p-d$ whether the seller sets her reserve price equal to $p-d$ or p, and seller j trades after setting price $p-d$.
2. the trading price is equal to $p-d$ (p) when the seller sets her reserve price equal to $p-d$ (p) and the seller fails to trade at price p[9].
3. the trading price is equal to p whether the seller sets her reserve price equal to $p-d$ or p, and seller j fails to trade after setting reserve price p.
4. the trading price is equal to $p-d$ (p) when the seller sets her reserve price equal to $p-d$ (p) and the seller trades at price p.

The seller gets a higher expected payoff by setting reserve price equal to $p-d$ if the effects of (1)-(3) outweigh the effect of (4). In the proof, we ignore the effect of (1) and (2) and show that the effect of (3) alone dominates (4). The effects of (1) and (2) are ignored because when $p-d=c$, the net surplus to seller j from trading at price $p-d$ is zero. Comparing (3) and (4), we can reinterpret our findings as follows: a seller posting price $p-d$ has higher chances of trading at price p than a seller posting price p. To understand the intuition behind this result, note that the expected number of sellers posting a reserve price equal to the trading price increases as the number of traders grows.

---

[9] It is easy to show that seller j will always trade after posting reserve price $p-d$ in this case

Competition from these sellers implies that with a significant probability a seller posting such reserve price fails to trade. Yet, if this seller posts a reserve price that is slightly lower than the trading price, she will trade with probability 1 and obtain a positive surplus.

Observe finally that the conclusion of the theorem is uniform in seller cost. It follows that for any distributions F and G, there must exist finite m and n such that there is a perfect Bayesian equilibrium in which all sellers set reserve prices equal to their true costs. We do not view this as the primary contribution of the paper and include the theorem more for completeness than anything else. Our demonstration that the perfect Bayesian equilibrium for the buyers bidding game makes the auction strategically equivalent to a seller's bid double auction from the sellers' point of view suggests a couple of arguments that already exist in the literature. (Rustichini, Satterthwaite, and Williams 1994) show that if a seller's bid double auction has a symmetric equilibrium it will converge asymptotically to an efficient auction. (Pesendorfer and Swinkels 2000) show asymptotic efficiency of double auctions in common value environments. Since private values represent a special case, their argument is suggestive about what should happen here. We get stronger results in the sense that fully efficient equilibria exist for finite numbers of buyers and sellers. This is due to the assumption that bids must lie in a grid[10]. Though this result is not particularly surprising given the asymptotic results, it could nonetheless be considered a contribution of this paper.

4 Conclusion

At least two remarks are in order about the results of this paper. First, the equilibrium in buyers' bidding game that we describe is not unique. Although the behavior that occurs in our equilibrium is plausible, alternative equilibria exist and do not generally guarantee efficient allocations. Examples in the paper illustrate this. We do not have a complete characterization of all equilibria in the dynamic bidding game, and of sellers' optimal behavior with respect to reserve prices when they believe that buyers are to play a different equilibrium.

Part of the job of an equilibrium in a decentralized market is to coordinate the matching decisions that buyers and sellers make. Coordination problems almost always have multiple equilibria and having to choose among them seems inevitable. Second-price auctions, for example, possess asymmetric equilibria in which bidders do not bid their true valuations. There are multiple equilibria in centralized mechanisms like double auctions as well.

Taking the multiplicity problem seriously, we do not want to suggest that traders will play the equilibrium described in this paper under all circumstances.

---

[10] This assumption is also what gets us around the impossibility result of Myerson and Satterthwaite. Dekel and Wolinsky is another paper that expoits the finiteness of the set of feasible bids. They show that buyers participating in a single good first price auction will bid their true valuations if the number of buyers is large. Their environment is considerably simpler than ours, so in fact they can go further. They show that bidding true valuation is a dominant strategy relative to a suitably constrained set of beliefs.

At the same time, the equilibrium behavior that we identify has a number of advantages. It is reasonable-looking, i.e. simple, requires very little computation on the part of traders, and is invariant to the form of the distributions from which costs and valuations are drawn. It also implements the efficient allocation. These properties make us believe that this equilibrium has focal nature, and eventually traders will learn to play it.

Of course, our model does not reproduce all details of the bidding behavior on the ebay, Amazon or other auction sites. These auctions typically possess additional dynamic aspects that we do not consider. For example, sellers enter at random times, as do buyers. Auctions close at different times. Furthermore, bidding is not completely costless on these sites. Roth and Ockenfels (2000) suggest network congestion and unexpected demands by the family as reasons why bidders may not be able to revise a bid as intended.

Nonetheless, our model does provide some insight into the workings of these institutions. In particular, bidding behavior of buyers in an auction on eBay cannot be determined in isolation. The choice to bid and the amount that one bids are largely determined by other alternative available at the time. In this sense, internet auctions are quite different from the one shot auctions that predominate in auction theory. This is a characteristic of the fact that internet auctions provide bidders with more information, and make it easier for them to communicate.

5 Appendix

Proof of Theorem 2:

At first, let us introduce the following notation. Let the *state of the game* be the array of buyers' valuations, sellers' standing bids together with the identities of buyers who have submitted them, the winning bids together with the identities of the winning bidders, the history of the standing bids and winning bids, plus the identity of the buyer who is the next to move, and the order in which buyers will move in the continuation. Typical state of the game will be denoted by Γ. Note that there is a one-to-one relationship between the nodes in the game and states of the game. Precisely, state of the game is a full description of the situation at the corresponding node in the game. Define *public state of the game* as the components of the *state of the game* that are publicly known. Specifically, public state of the game includes the standing bids, the winning bids, the identities of the winning bidders, the history of all these, and the order of moves.

We will say that state of the game Γ is *regular* if every buyer's valuation is at least as large as any high bid that he holds at Γ. Note that whether the state of the game is regular or not is unobservable.

We will say that bidder i's position is *consistent with* σ^* if each of i's high bids has the following property: if i is a high bidder with seller j then either (i) i's bid is one grid point above seller j's standing bid and no seller has a lower standing bid than does seller j; or (ii) i's high bid is equal to j's standing bid. It is not hard to see from the definition that on the path generated by σ^* all states of the game are regular, and that every buyer's position in these information sets will be consistent with σ^*.

If in state Γ one or more buyers are high bidders at more than one seller, we will replace the continuation of the true game from the node that corresponds to Γ with the one where additional 'phantom' buyers are added to the game, but where each buyer (real or phantom one) has a winning bid at one seller at most. The phantom buyers are added in the following way. If buyer i is a winning bidder at multiple (say, $l > 1$) sellers, choose any one of i's highest winning bids and consider that buyer i possesses only this winning bid. For each of i's $l-1$ other high bids, create a phantom buyer i^k $(2 \leq k \leq l-1)$ whose valuation is equal to this winning bid. Then, define $\tilde{\mathbf{v}}$ to be the vector of the valuations of the real and phantom bidders, and the current standing bids in state Γ. If there are $\tilde{m} \geq m$ buyers and phantom buyers, then $\tilde{\mathbf{v}}$ has dimension $\tilde{m}+n$. Let $\tilde{v}_r$ is the r-th smallest element of vector $\tilde{\mathbf{v}}$.

Define B_Γ to be the *smallest* set of buyers (including phantom buyers) at state Γ such that if B_Γ contains m_Γ bidders, every buyer who is not in B_Γ is high bidder with a seller whose standing bid strictly exceeds $\tilde{v}_{m_\Gamma}$.

Lemma 5 *If state of the game Γ is regular, then the set B_Γ is non-empty and unique. Furthermore if B_Γ contains m_Γ buyers, no buyer in B_Γ is high bidder at a seller with standing bid above $\tilde{v}_{m_\Gamma}$.*

Proof. Let B_0 be the set of buyers who are not high bidders at any seller. Suppose that B_0 includes τ_0 buyers. By definition, $B_0 \subset B_\Gamma$. Let s_0 be the lowest standing bid. If $\tilde{v}_{\tau_0} < s_0$ then $B_\Gamma = B_0$, and we are done. Otherwise, B_Γ must contain *all* buyers who are high bidders at sellers with standing bid s_0. To see this, consider any candidate set B_1 with $\tau_1 \geq \tau_0$ members which contains B_0. Since $\tilde{v}_{\tau_1} \geq \tilde{v}_{\tau_0} \geq s_0$, a high bidder at s_0 does not satisfy the condition for exclusion from B_1.

We can now apply the same argument recursively. Let B_l be the set that includes all buyers from B_0 and all buyers who are high bidders at the lowest, second lowest, ..., l-th lowest standing bids. Suppose that there are τ_l buyers in B_l, and no subset of B_l satisfies the definition of B_Γ. If $\tilde{v}_{\tau_l} < s_{l+1}$, where s_{l+1} is the $l+1$-th lowest standing bid, then $B_\Gamma = B_l$. Otherwise, B_Γ must contain *all* buyers who are high bidders with sellers whose standing bids are less than or equal to s_{l+1}. Continue in this way until l' s.t. $B_{l'}$ satisfies the definition of B_Γ or until all buyers have been included in B_Γ.

It is immediate that no buyer in B_Γ is a high bidder at a seller whose standing bid exceeds $\tilde{v}_{m_\Gamma}$ (where m_Γ is the number of bidders in Γ), for such a bidder can be excluded from B_Γ to form a strictly smaller set satisfying the definition. ■

Figure 3 provides an example of a situation with phantom bidders. Buyers' valuations b_1, b_2, b_3 and sellers' reserve prices s_1, s_2 are the same as in Figure 1. However, in the state which we are considering buyer 1 has submitted two high bids, p_1 with seller 1 and $p_2 = b_3$ with seller 2, but no other buyer has entered

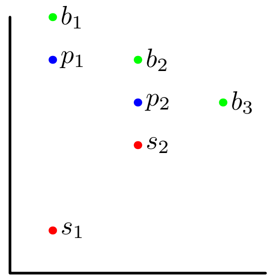

Figure 3: Phantom Bidders

yet. The standing bids remain s_1 at seller 1 and s_2 at seller 2. Since buyer 1 is high bidder at two sellers, we generate a phantom buyer whose valuation and bid are both equal to p_2. There are then 4 buyers including the phantom buyer, instead of only 3 in the original game). The vector $\tilde{\mathbf{v}}$ of valuations and standing bids in increasing order is given by $(s_1, s_2, p_2, b_3, b_2, p_1, b_1)$. B_Γ is the whole set of buyers, including the phantom buyer. Consequently, $m_\Gamma = 4$ and $\tilde{v}_4 = p_2 = b_3$. None of the buyers is a high bidder with a seller whose standing bid is above p_2, so this is the right choice for B_Γ.

To understand the outcome on the continuation path where all buyers follow σ^*, assume that buyer 3 enters first. Buyer 3 will bid up the standing bid with each of the two sellers until both standing bids reach his valuation b_3. Buyer two will then choose between the two sellers randomly. Whichever choice buyer 2 makes, he will trade at price p_2, while buyer 1 will trade either at price p_2 or p_1 depending on where buyer 2 ends up submitting his bid first.

Figure 3 illustrates why σ^* may not generate uniform prices in the continuation after arbitrary states. It also shows why buyers will want to avoid submitting high bids early, even though each seller runs a second price auction. Buyer 1 may have to pay p_1 instead of p_2 if bidder 2 bids at seller 1 first.

To see another example where non-uniform prices can result, suppose that the high bids p_1 and p_2 are held by buyers 1 and 2 respectively and eliminate buyer 3 from the data in Figure 3. In this case no buyer is a high bidder at more than one seller, but both buyers use strategies that are inconsistent with σ^*. If they subsequently revert to σ^* then bidding will stop and buyer 1 will trade at the current standing bid s_1, and buyer 2 will trade at price s_2.

Lemma 6 *Consider any regular state Γ and suppose that all buyers use σ^* in the continuation. Then no buyer whose position is consistent with σ^* from the start of the game up to Γ will trade at a price above $\tilde{v}_{m_\Gamma}$.*

Proof. The proof is by contradiction. Thus, suppose that some buyer i from B_Γ whose position is consistent with σ^* in all states of the game up to Γ ends up trading at price $p > \tilde{v}_{m_\Gamma}$ with some seller j. We need to consider three cases:

Case (a). When buyer i submits her last bid $b^i > \tilde{v}_{m_\Gamma}$, the standing bid at seller j is equal to p. Since i follows σ^*, the lowest standing bid at this point must be p.

Let k be the number of buyers and phantom bidders in B_Γ whose valuations are greater than $\tilde{v}_{m_\Gamma}$. By the definition of $\tilde{v}_{m_\Gamma}$, there must be at least k sellers whose standing bids in state Γ are $\tilde{v}_{m_\Gamma}$ or lower. All these sellers have standing bids of at least p when i submits bid b^i at seller j. When buyer i submits her bid b^i, he cannot be a high bidder with some seller, so there are only $k-1$ bidders and phantom bidders who can be high bidders with these k sellers at standing bids of at least p. So at least one of the bidders and phantom bidders other than i must be high bidder with more than one seller. Since by construction each bidder and phantom bidder is high bidder with at most one seller in state Γ, some buyer must have submitted multiple high bids above $\tilde{v}_{m_\Gamma}$ in the continuation following Γ. This is inconsistent with part (i) in the definition of σ^*.

Case (b). Buyer i submits his last bid b^i in the continuation following state Γ when the standing bid s_j at seller j is below p.

Since i follows σ^*, $b^i = s_j + d$. Hence, $b^i = p$. As buyer i ends up trading at p, bid b^i makes him the winning bidder. However, the standing bid cannot change at this point, because if it did, it would imply that $b^i > s^j + d$.

Since the standing with seller j eventually rises to p, some other buyer, say i', must submit a bid p with seller j after i has submitted her last bid $b^i = p$ and while the standing bid is still s^j. The identity of the winning bidder at seller j has changed after the standing bid has reached s^j. Therefore, buyer i' will bid at j only if s^j is the lowest standing bid and the identity of the winning bidder has changed at all sellers with standing bid s_j including sellers whose standing bids are below $\tilde{v}_{m_\Gamma}$ in state Γ. This implies that all k sellers who in state Γ have standing bids below $\tilde{v}_{m_\Gamma}$ have winning bids above $s^j \geq \tilde{v}_{m_\Gamma}$.

Since the state Γ is regular and since buyer i' will not submit bid p if he is already high bidder with some seller, there are only $k-1$ buyers from B_Γ and phantom bidders who can be high bidders with these k sellers at standing bids of at least $p = s_j + d$. So at least one of the bidders and phantom bidders other than i' must be high bidder with more than one seller. By construction, each bidder and phantom bidder is high bidder with at most one seller in state Γ. Therefore, some buyer must have submitted multiple high bids above $\tilde{v}_{m_\Gamma}$ in the continuation following Γ. This is inconsistent with part (i) in the definition of σ^*.

Case (c). At state Γ i is holding her final winning bid $b^i > \tilde{v}_{tau}$ at seller j with standing bid $s^{j'} \leq \tilde{v}_{m_\Gamma}$. Since i is in B_Γ and his position is consistent with σ^*, we must have $s^{j'} = \tilde{v}_{tau}$ and $p = b^i = \tilde{v}_{tau} + d$. Since i's position is consistent with σ^* on the whole path to Γ, bidder i's last bid has changed the identity of the winning bidder but not the standing bid.

Since the standing with seller j eventually rises to p, some other buyer, say i', must submit a bid p with seller j after i has submitted her last bid $b^i = p$ and while the standing bid is still s^j.

Since the standing with seller j eventually rises to $p = \tilde{v}_{tau} + d$, some other buyer, say i'', must submit a bid $\tilde{v}_{tau}$ with seller j in the continuation following state Γ. To complete the proof, apply the same argument as in (b). ■

Lemma 7 *Let* Γ *be a regular state. If all buyers use* σ^* *in the continuation, then no trader will trade at a price below* v_{m_Γ}.

Proof. Suppose that some buyer trades at price $p < v_{m_\Gamma}$. Then there is at least one seller whose standing bid in state Γ is strictly less than v_{m_Γ}. Let the number of such sellers be r_1. Also, let B_1 be the set of buyers with valuations equal to or greater than v_{m_Γ} and who in state Γ are either winning bidders with sellers whose standing bids are below v_{m_Γ} or are not winning bidders at any seller. Denote the number of buyers in B_1 by k_1. By definition $B_1 \subset B_\Gamma$.

First, we will show that $k_1 > r_1$. The proof is by contradiction. Thus, suppose otherwise i.e. $k_1 \leq r_1$. By definition of v_{m_Γ}, there are $k_2 > r_1$ buyers and phantom buyers in B_Γ whose valuations are v_{m_Γ} or higher. Therefore, there are $k_2 - k_1 > 0$ buyers in B_Γ each of whom has valuations at or above v_{m_Γ} and is a high bidder with a seller whose standing bid in Γ is equal to v_{m_Γ}. Eliminating these buyers from set B_Γ we obtain a strictly smaller set B'_Γ.

Thus, B'_Γ is the union of the set of buyers who have valuations below v_{m_Γ} (let there be k_0 of them) and set B_1. The total number of buyers in B_Γ is $m_{\Gamma'} = k_0 + K_1$. Since $r_1 > k_1$, $v_{m_{\Gamma'}} < v_{m_\Gamma}$. Hence, every buyer who is not in $B_{\Gamma'}$ is a high bidder with a seller whose standing bid exceeds $v_{m_{\Gamma'}}$. This contradicts the fact that B_Γ is the smallest set that has such property in Γ.

If trade occurs in the continuation at price $p < v_{m_\Gamma}$, then on the continuation path no buyer submits a bid when the standing bid with some seller is p. Since $k_1 > r_1$, it follows that in the continuation following Γ at least one of the buyers from the set B must has become a high bidder with a seller whose standing bid in Γ is v_{m_Γ} or higher. Submitting such a bid when there is a standing bid below v_{m_Γ} is inconsistent with σ^*, so at least one buyer has to deviate from σ^* in order for this outcome to occur. This contradiction proves the result. ■

Lemma 8 *Let* Γ *be any regular state. If all buyers use* σ^* *in the continuation, then every buyer who is not in* B_Γ *trades with the seller with whom he is high bidder in state* Γ *at this seller's standing bid in* Γ.

Proof. Buyers who are not in B_Γ are all high bidders with sellers whose standing bids exceed v_{m_Γ}. If they all use σ^* in the continuation, none of them will be displaced as a higher bidder unless some buyer from B_Γ submits a bid with one of the sellers whose standing bid is strictly above v_{m_Γ}. Such a bid cannot occur because by Lemma 7, a buyer from B_Γ whose valuation is above v_{m_Γ} and who follows σ^* is guaranteed to trade if he submits a bid equal to $v_{m_\Gamma} + d$ at one of the sellers whose standing bid in state Γ does not exceed v_{m_Γ}. ■

One implication of Lemmas 7- 8 is that any buyer with valuations above v_{m_Γ} whose initial position is consistent with σ^* and who follows the strategy σ^* in the continuation will trade at a price equal to v_{m_Γ} provided that all other buyers follow σ^*. Consider the implications of this result starting from the initial state Γ_0 in which no buyer has yet submitted a bid. In this state , all buyers' positions are consistent with σ^* (trivially since no buyer is a high bidder). Therefore, if all buyers follow σ^* in the continuation, then all trades will take place at price v_m.

Lemma 9 *Let Γ be a regular state, and suppose that every buyer's position is consistent with σ^*. Then in the continuation no buyer can increase his payoff by choosing an action other than that specified by σ^* provided that the other buyers follow σ^*.*

Proof. By lemmas 7 and 6, if in state Γ every buyer's position is consistent with σ^* and all buyers follow σ^* then: (i) all trades will take place at price v_{m_Γ}. (ii) A buyer i whose valuations strictly exceeds v_{m_Γ} will trade for sure at this price.

Consider now any continuation path induced by a series of deviations by buyer i when the other buyers follow σ^*. Each such path has finite length since the other bidders will eventually stop bidding and i cannot bid against himself. Thus i must trivially revert to σ^* at some point along the path (for example when he stops bidding). Since i must always submit bids at least as high as an existing standing bid, the lower bound on i's trading price given by Lemma 7 along any such path can never be lower than v_{m_Γ}. $\blacksquare$

Say that buyers' beliefs are σ^*-*consistent* if each buyer believes that: (i) every other buyer's valuation is drawn from the prior distribution of valuations conditional on each high bidder's valuation being at least as large as the maximum of her high bids; (ii) each of the other buyers' high bids is consistent with σ^*. In other words, buyers' beliefs are σ^*-consistent if every buyer believes that the only states that occur with positive probability are regular states in which all other buyers' positions are consistent with σ^*.

The proof that σ^* constitutes a Bayesian equilibrium in the bidding game given sellers' announced reserve prices then follows immediately from Lemma 9 since deviations from σ^* may raise, but can never lower a buyer's trading price in every information set that the buyer thinks occurs with positive probability.

Proof of corollary 3:

Suppose that seller j posts reservation price s s.t. $s < v_m$. According to theorem 2, when buyers use strategy σ^* all trades occur at price v_m. Consider the last bid in the bidding game submitted by some buyer i'. It must be equal to v_m or $v_m + d$. If the last bid is equal to $v_m + d$, then the lowest standing bid at this point must be v_m. Hence, seller j will trade.

Suppose now that the last bid in the game is equal to v_m and is submitted at some seller j'. Then before the last bid is submitted the standing bid at j' must be equal to $v_m - d$, and after the last bid is submitted the standing bid at j' must increase to v_m. Therefore, bidder i is not a winning bidder at the end of the auction. This can happen only if at this point the standing bid at seller j is v_m. Hence, seller j must trade in this case also.

Consider set of sellers N_2 who post reserve prices equal to v_m. A seller from N_2 trades only if a buyer from M_3 bids with her. After accounting for sellers from N_1 who trade for sure, the number of sellers from M_3 who are available to bid with sellers from N_2 is at least $m_3 - n_1$. This gives the lower bound on the number of sellers from N_2 who trade.

To obtain the upper bound, note that on the equilibrium path buyers from M_2 will bid only with sellers from N_1 while the standing bids at these sellers are below v_m. Consider the first time t when the lowest standing bid in the market reaches v_m. With a positive probability, the realizations of random order of bidding and the randomizations by buyers between the sellers among whom they are indifferent is such that at time t m' buyers from M_2 are the high bidders at sellers from N_1, where m' is between $\max\{0, n_1 - m_3\}$ and $\min\{m_2, n_1\}$. Also, with a positive probability all buyers from M_3 who at time t are not winners yet, will bid at sellers from N_2 first. The number of buyers from M_3 who will bid in this way is equal to: $m_3 - (n_1 - m')$. Substituting for m' we get the upper bound on the number of sellers from N_2 who trade.

The proof establishing the lower and upper bounds on the number of buyers from M_2 who trade is similar and is, therefore, omitted.

Proof of theorem 4:

Consider seller z with cost c. We will demonstrate that seller z's expected payoff decreases in her reserve price p if $p > c$. To accomplish this, we will compare the expected payoffs that the seller gets when she posts reserve prices equal to p and $p - d$.

In the previous section we have demonstrated that all trades will be completed at a uniform trading price equal to v_m-the m-th lowest element in $\mathbf{v}$, the vector of the true buyers' valuations and sellers' reserve prices which, by assumption of the theorem, are equal to the true costs for all sellers except z. Thus, the trading price will be equal to p^T if and only if the following two necessary and sufficient conditions hold. First, the number of sellers and buyers whose reserve prices and valuations respectively are strictly below p^T does not exceed $m - 1$. Second, the number of sellers and buyers whose reserve prices and valuations respectively are no greater than p^T is at least m.

We will use the following notation that has been introduced above: $m_1/m_2/m_3$ is the number of buyers with valuations *strictly below p / equal to p /strictly above p.* Similarly, n_1/n_3 is the number of sellers with costs *strictly below p/strictly above p.* Also, let n_2' be the number of sellers, other than z with costs equal to p. Obviously, $n_2' = n_2 - 1$, $m_1 + m_2 + m_3 = m$ and $n_1 + n_2' + n_3 = n - 1$.

At first, let us establish the following two claims.

Claim 1. *Suppose that if seller z sets reserve price p, then the trading price is p_T s.t.$p < p_T$. Then, if seller z sets a different reserve price $p' < p_T$, the trading price will also be p_T.* The buyer will trade in both cases.

Proof: The trading price is equal to v_m which is not affected by a change in the reserve price p set by z as long as $p < v_m$.

Claim 2. *Suppose that if seller z sets reserve price p, then the trading price is p_T s.t. $p > p_T$. Then, if seller z sets reserve price $p'' > p_T$ the trading price will also be p_T.* The buyer will fail to trade in both cases.

Proof: The trading price is equal to v_m which is not affected by a change in the reserve price p set by z as long as $p > v_m$.

Say that price p is pivotal if the trading price is equal to p when seller z sets her reserve price equal to p. Claims 1 and 2 imply that seller z's payoffs from setting reserve price equal to $p - d$ or p may be different only if at least one of these reserve prices is pivotal.

Let $P(\Omega)$ denote the probability of event Ω and $E(y)$ $E(y|\Omega)$ denote the expectation (conditional expectation given event Ω) of the random variable y. Then the following lemma presents sufficient condition for theorem 4 to hold.

Lemma 10 *Seller z with cost c gets a higher expected payoff by setting reserve price $p - d$ rather than p for $p > c$ if the following condition holds:*

$$\begin{aligned} &P(p \text{ is pivotal }, p-d \text{ is not pivotal, seller posting } p \text{ fails to trade }) \geq \\ &P(p \text{ is pivotal }, p-d \text{ is pivotal, seller posting } p \text{ trades }) \end{aligned} \tag{1}$$

Proof. By Claims 1 and 2, it is sufficient to compare the seller's expected payoffs from setting her reserve price equal to $p - d$ and p when at least one of these prices is pivotal. Let us consider all such cases. Note that a seller may fail to trade when her reserve price is pivotal, because there may be other sellers who post this reservation price. Alternatively, there may be buyers with valuations equal to the pivotal price.

1. If p is pivotal, but $p-d$ is not, then irrespective of seller z's reserve price, v_m and hence the trading price are equal to p. This follows because in this case we have: $m_1 + n_1 < m - 1 \leq m_1 + n_1 + m_2 + n_2'$. Consequently, if z sets reserve price $p - d$ she trades at price p for sure. If she sets reserve price p, she may fail to trade if $m_1 + n_1 + m_2 + n_2' \geq m$.

2. If $p-d$ is pivotal, but p is not, then irrespective of seller z's reserve price, v_m (and hence the trading price) is equal to $p - d$. This follows because in this case we must have $m_1 + n_1 \geq m$. Consequently, if z sets reserve price p she fails to trade. If she sets reserve price $p - d$, she may or may not trade.

3. If both $p - d$ and p are pivotal, we must have: $m_1 + n_1 = m - 1$. Then if seller z posts reserve price $p-d$ she will trade at this price for sure. If the seller sets reserve price p, the trading price will be equal to p but seller j may fail to trade if $m_2 = n_2 > 0$.

Summing up these effects we conclude that seller z with cost $c < p$ obtains a higher payoff by setting reserve price $p - d$ rather than p if and only if the following inequality holds:

$$\begin{aligned}
&(p-c)P(p \text{ is pivotal, } p-d \text{ is not pivotal, seller posting } p \text{ fails to trade })+\\
&(p-d-c)P(p-d \text{ is pivotal, } p \text{ is not pivotal, seller posting } p-d \text{ trades })+\\
&(p-d-c)P(p-d \text{ is pivotal, } p \text{ is pivotal })\geq\\
&(p-c)P(p \text{ is pivotal, } p-d \text{ is pivotal, seller posting } p \text{ trades }) \qquad (2)
\end{aligned}$$

Obviously, (2) holds when (1) holds, and these two inequalities are equivalent when $c = p - d$. ■

Lemma 11 *Inequality (1) holds if:*

$$\begin{aligned}
&E\left(\frac{1+n_2'+n_1+m_1-m}{n_2'+1}|m_1+n_1<m-1\leq m_1+n_1+m_2+n_2'\right)\times\\
&\times P(m_1+n_1<m-1\leq m_1+n_1+m_2+n_2')\\
&\geq E(\frac{1}{n_2'+1}|m_1+n_1=m-1)P(m_1+n_1=m-1) \qquad (3)
\end{aligned}$$

Proof. First of all, p is pivotal and $p-d$ is not pivotal if and only if $m_1+n_1<m-1\leq m_1+n_1+n_2'+m_2$. Also, p and $p-d$ are both pivotal if and only if $m_1+n_1=m-1$.

Next, let us compute the upper bound on the probability with which seller z posting pivotal price p trades conditional on the number of buyers and sellers in each category. By rule (a) of σ^*, seller z can trade only with one of the m_3 buyers whose valuations are above p. By corollary 3, n_1 sellers who post reserve prices below p trade for sure. Some of these n_1 sellers may trade with buyers whose valuations are equal to p. Therefore, the number of buyers who trade with sellers posting p is at most $m_3+\min\{0,m_2-n_1\}$, and will be lower if some buyers with valuations equal to p do not trade and some buyers with valuations above p trade with sellers whose reserve prices are below p.

Seller z competes with the other n_2' sellers who post reserve price p. Since all buyers follow strategy σ^*, a buyer who chooses among such sellers will randomize between them with equal probability. Therefore, the conditional probability that seller z trades is at most

$$P^t \equiv \min\{1, \frac{m_3+\min\{0,m_2-n_1\}}{n_2'+1}\}$$

Correspondingly, the lower bound on the probability that seller z fails to trade conditional on the number of buyers and sellers in each category is

$$P^f \equiv \max\{0, \frac{n_2'+1-m_3-\min\{0,m_2-n_1\}}{n_2'+1}\}$$

Finally, taking expectation of P^f conditional on the event that both $p-d$ and p are pivotal, which is equivalent to $m_1+n_1<m-1\leq m_1+n_1+m_2+n_2'$, we

obtain that the left-hand side of (1) is at least as large as the left-hand side of (3). Taking expectation of P^t conditional on $m_1 + n_1 = m - 1$ and simplifying we obtain that the right-hand side of (1) is no greater than the right-hand side of (3). ■

To complete the proof of the theorem it remains to show that (3) holds when m and n are sufficiently large. First of all, note that the right-hand side of (3) is equal to $A_1(\alpha) + A_2(\alpha)$ where

$$A_1(\alpha) \equiv$$

$$\sum_{\hat{m}_1=[\alpha m]+1}^{m-1} E\left(\frac{1}{n_2'+1}|n_1 = m-1-\hat{m}_1\right) P(m_1 = \hat{m}_1)P(n_1 = m-1-\hat{m}_1)$$

and

$$A_2(\alpha) \equiv$$

$$\sum_{\hat{m}_1=\max\{0,m-n\}}^{[\alpha m]} E\left(\frac{1}{n_2'+1}|n_1 = m-1-\hat{m}_1\right) P(m_1 = \hat{m}_1)P(n_1 = m-1-\hat{m}_1)$$

where $\alpha \in [0,1]$ and $[\alpha m]$ is defined as the largest integer not exceeding αm.

The expression on the left-hand side of (3) is greater than $B_1 + B_2$ where B_1 and B_2 are defined as follows:

$$B_1 \equiv$$

$$\sum_{\hat{m}_1=\max\{0,m-n\}}^{m-2} E\left(\frac{1+n_2'+n_1+m_1-m}{n_2'+1}|m_1 = \hat{m}_1+1, n_1 = m-\hat{m}_1-3\right) \times$$

$$\times P(m_1 = \hat{m}_1+1)P(n_1 = m-\hat{m}_1-3)$$

and

$$B_2 \equiv$$

$$\sum_{\hat{m}_1=\max\{2,m-n+4\}}^{m-1} E\left(\frac{1+n_2'+n_1+m_1-m}{n_2'+1}|m_1 = \hat{m}_1-2, n_1 = m-\hat{m}_1-1, n_2' \geq 4\right) \times$$

$$\times P(m_1 = \hat{m}_1-2)P(n_1 = m-\hat{m}_1-1, n_2' \geq 4)$$

To complete the proof of the theorem, let us prove the following lemma:

Lemma 12 *$\exists \alpha \in (0,1)$ s.t. if and m (and hence n) is large enough then $A_1(\alpha) < B_1$ and $A_2(\alpha) < B_2$.*

Proof. The following expressions are used extensively in the sequel. Let $a+b+c = n-1$ and $r+s+t = m$. Then we have:

$$P(n_1 = a, n_2 = b, n_3 = c) = \frac{(n-1)!}{a!b!c!} G(p-1)^a g(p)^b (1-G(p))^c \quad (4)$$

$$P(m_1 = r, m_2 = s, m_3 = t) = \frac{m!}{r!s!t!} F(p-1)^r f(p)^s (1-F(p))^t \quad (5)$$

Part 1: $B_1 > A_1(\alpha)$.

Step 1. Using (5), we have for $2 \leq \hat{m}_1 \leq m-1$:

$$P(m_1 = \hat{m}_1 - 2) = \frac{(1-F(p-d))^2}{(F(p-d))^2} \frac{\hat{m}_1(\hat{m}_1 - 1)}{(m-\hat{m}_1)(m-\hat{m}+1)} P(m_1 = \hat{m}_1) \quad (6)$$

Since $1 - F(p-d) > \bar{f} > 0$, $\exists m'$ and $\alpha_1 > 0$ s.t. if $\hat{m}_1 > \alpha_1 m'$, $P(m_1 = \hat{m}_1 - 2) > P(m_1 = \hat{m}_1)$. ■

Step 2. When $n_1 = m - \hat{m} - 1$, then $n_2' + n_3 = \hat{m}_1 + n - m = \hat{m}_1 - \frac{k-1}{k} m$. Let $\alpha = \max\{\alpha_1, \frac{k-1/2}{k}\}$. (Note that $\alpha < 1$.) In this case, $\hat{m}_1 \geq \alpha m$ implies that $n_2 + n_3 \geq \frac{m}{2k}$. Since $g(p) \geq \bar{g} > 0$, the law of large numbers implies that there is m'' s.t. if $m > m''$ then $P(n_2' \geq 4 | n_2' + n_3 \geq \frac{m}{2k}) > 5/6$, and hence $P(n_1 = m - \hat{m}_1 - 1, n_2' \geq 4) > \frac{5}{6} P(n_1 = m - \hat{m}_1 - 1)$. Therefore,

$$E(\frac{1}{n_2'+1} | n_1 = m - \hat{m}_1 - 1) P(n_1 = m - \hat{m}_1 - 1) =$$

$$E(\frac{1}{n_2'+1} | n_2' < 4, n_1 = m - \hat{m}_1 - 1) P(n_2' < 4, n_1 = m - \hat{m}_1 - 1) +$$

$$E(\frac{1}{n_2'+1} | n_2' \geq 4, n_1 = m - \hat{m}_1 - 1) P(n_2' \geq 4, n_1 = m - \hat{m}_1 - 1)$$

$$< (1 \times 1/6 + 1/5 \times 5/6)\, P(n_1 = m - \hat{m}_1 - 1) = 1/3 P(n_1 = m - \hat{m}_1 - 1) \quad (7)$$

On the other hand, we have:

$$E\left(\frac{1 + n_2' + n_1 + m_1 - m}{n_2' + 1} | m_1 = \hat{m}_1 - 2, n_1 = m - \hat{m}_1 - 1, n_2' \geq 4\right) \times$$

$$\times P(n_1 = m - \hat{m}_1 - 1, n_2' \geq 4)$$

$$\geq \frac{2}{5} P(n_1 = m - \hat{m}_1 - 1, n_2' \geq 4) >$$

$$\frac{1}{3}P(n_1 = m - \hat{m}_1 - 1) \tag{8}$$

Step 3. Choose $m = \max\{m_1, m_2\}$. Combining (6), (7) and (8), we conclude that the term corresponding to each $\hat{m}_1 \geq \alpha m$ in $A_1(\alpha)$ is dominated by the corresponding term in B_1. Therefore, $B_1 > A_1(\alpha)$.

Part 2. $B_2 > A_2(\alpha)$.

Proof. Consider α defined in part 1. Fix $\hat{m}_1$ s.t. $\max\{0, m-n\} \leq \hat{m}_1 \leq [\alpha m]$, $\hat{n}_1 = m - \hat{m}_1 - 1$ and $\hat{n}_2$ s.t. $0 \leq \hat{n}_2 \leq n - 1 - \hat{n}_1$. Let us show that the following inequality holds for sufficiently large m. (Note that $\hat{n}_1 > 2$ when m is not too small).

$$E\left(\frac{n_2+1+n_1+m_1-m}{n_2+1} | n_1 = \hat{n}_1 - 2, m_1 = \hat{m}_1 + 1, n_2 = \hat{n}_2 + 2\right) \times$$

$$P\left(n_1 = \hat{n}_1 - 2, m_1 = \hat{m}_1 + 1, n_2 = \hat{n}_2 + 2\right) >$$

$$\frac{1}{n_2+1}P(n_1 = \hat{n}_1, m_1 = \hat{m}_1, n_2 = \hat{n}_2) \tag{9}$$

To establish this, note the following sequence of equalities which holds by computation:

$$E\left(\frac{n_2+1+n_1+m_1-m}{n_2+1} | n_1 = \hat{n}_1 - 2, m_1 = \hat{m}_1 + 1, n_2 = \hat{n}_2 + 2\right) \times$$

$$P\left(n_1 = \hat{n}_1 - 2, m_1 = \hat{m}_1 + 1, n_2 = \hat{n}_2 + 2\right) =$$

$$\frac{\hat{n}_2+1}{\hat{n}_2+3}P\left(n_1 = \hat{n}_1 - 2, m_1 = \hat{m}_1 + 1, n_2 = \hat{n}_2 + 2\right) =$$

$$\frac{\hat{n}_2+1}{\hat{n}_2+3}P(n_1 = \hat{n}_1, m_1 = \hat{m}_1, n_2 = \hat{n}_2)\times$$

$$\frac{g(p)^2}{G(p-1)^2}\frac{(\hat{n}_1-1)\hat{n}_1}{(\hat{n}_2+2)(\hat{n}_2+1)}\frac{F(p-1)}{1-F(p-1)}\frac{m-\hat{m}_1}{\hat{m}_1+1} =$$

$$\frac{1}{\hat{n}_2+1}P(n_1 = \hat{n}_1, m_1 = \hat{m}_1, n_2 = \hat{n}_2)\times$$

$$\frac{g(p)^2}{G(p-1)^2}\frac{F(p-1)}{1-F(p-1)}\frac{(\hat{n}_2+1)(\hat{n}_1-1)\hat{n}_1}{(\hat{n}_2+2)(\hat{n}_2+3)}\frac{m-\hat{m}_1}{\hat{m}_1+1}$$

Since $\hat{m}_1 \leq \alpha m$ and $\hat{n}_1 = m - 1 - \hat{m}_1$, it follows that $m - \hat{m}_1 \geq (1-\alpha)m$ and $\hat{n}_1 \geq (1-\alpha)m - 1$. Also, $\hat{n}_2 \leq n - 1 - \hat{n}_1$ implies that $\hat{n}_2 \leq \left(\frac{1}{k} + \alpha - 1\right) m$. Therefore,

$$\frac{(\hat{n}_2+1)(\hat{n}_1-1)\hat{n}_1}{(\hat{n}_2+2)(\hat{n}_2+3)} \frac{m-\hat{m}_1}{\hat{m}_1+1} \geq$$

$$\frac{1}{2} \frac{((1-\alpha)m-2)\,((1-\alpha)m-1)}{\left(\frac{1}{k}+\alpha-1\right)m+3} \frac{(1-\alpha)m}{\alpha m+1}$$

The right-hand side of the above inequality if of the same order as m. Therefore, (9) holds when m is sufficiently large (leaving α unchanged).

Finally, note that (9) implies that the term corresponding to each $\hat{m}_1 \geq \alpha m$ in $A_2(\alpha)$ is dominated by the corresponding term in B_2. Therefore, $B_2 > A_2(\alpha)$ holds. ■

References

Ausubel, L. M., and P. Milgrom (2001): "Ascending Auctions with Package Bidding," University of Maryland Working paper.

Bajari, P., and A. Hortacsu (January, 2000): "Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," Stanford Working paper http://home.uchicago.edu/ hortacsu/ebay$_f$*inal.pdf*.

Carare, O. (2001): "Need for Speed: Demand Estimation Using Auction Data," Rutgers University working paper http://www.rci.rutgers.edu/ carare/jmkt/ns.pdf.

Krishna, V., and M. Perry (1998): "Efficient Mechanism Design," Pennsylvania State University manuscript.

McAfee, P. (1993): "Mechanism Design by Competing Sellers," *Econometrica*, 61(6), 1281–1312.

Morneau, J. (2001): "E-Marketplaces Start to Take Off," *TechWeb News, http://www.techweb.com.*

Pesendorfer, W., and J. Swinkels (2000): "Efficiency and Information Aggregation in Auctions," *American Economic Review.*

Peters, M. (1997): "On the Equivalence of Walrasian and Non-Walrasian Equilibria in Contract Markets," *Review of Economic Studies*, 64(2), 241–265.

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

Roth, A., and A. Ockenfels (2000): "Last Minute Bidding and the Rules for Ending Second Price Auctions: Theory and Evidence from a Natural Experiment on the Internet," Discussion paper, Harvard University.

Roth, A. E., and M. A. Sotomayor (1990): *Two-Sided Matching - A Study in Game Theoretic Modelling and Analysis*, Econometric Society Monographs. Cambridge University Press, Cambridge MA.

Rustichini, A., M. A. Satterthwaite, and S. R. Williams (1994): "Convergence to Efficiency in a Simple Market with Incomplete Information," *Econometrica*, 62(5), 1041–1063.

Satterthwaite, M. A., and S. Williams (1089): "The Rate of Convergence to Efficiency in the Buyer's bid Double Auction as the Market Becomes Large," 56, 477–498.

Wilson, R. (1985): "Incentive Efficiency of Double Auctions," 53, 1101–1115.

Wilson, T., and T. Mullen (2000): "E-Business Exchanges Fight For Survival," *InternetWeek, http://www.internetwk.com.*

Zheng, M. (2001): "Bidding Behavior on Ebay," Discussion paper, University of Toronto, http://www.chass.utoronto.ca/ mzheng/papers/eBay.pdf.