

Internet Auctions with Many Traders

Michael Peters and Sergei Severinov *

First Version Feb/01 This version July 3, 2001

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 decentralized 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.

1 Introduction

One of the central features of auction theory is the centralized nature of trade. Specifically, the mechanism designer, or auctioneer, establishes trading rules, then processes information contained in messages sent by traders, and adjusts prices. Such organization of trade turns out to be quite successful, as efficient outcomes can in many cases be implemented via an auction. Gul and Stacchetti (2000) and Ausubel (2000) have recently demonstrated that auctions can achieve efficiency even when multiple and heterogeneous objects are sold.

Less is known about multilateral trading environments that involve many buyers and sellers. The analysis of such environments in the literature has also relied on the model of centralized processing of demand and supply information (see, for example, Wilson (1985), Satterthwaite and Williams (1989a), Satterthwaite and Williams (1989b), Rustichini, Satterthwaite, and Williams (1994), Herings (1997)). To some extent, this line of research can be seen as a continuation of a standard competitive model which uses a fictitious auctioneer who collects demand and supply information, then uses some sort of tâtonnement process to compute market clearing prices.

One of the most powerful results in the theory of centralized multilateral trading concerns the incentive efficiency of double auctions (see Wilson (1985)). Rustichini, Satterthwaite, and Williams (1994), following previous work by Satterthwaite and Williams (1989a), show that surplus that is lost in centralized double auctions shrinks to zero very

*Department of Economics, University of Toronto, 150 St. George St, Toronto, ON M5S 3G7, Canada; email: peters@economics.utoronto.ca, and Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706; email: sseverin@ssc.wisc.edu.

quickly as the number of traders gets large. Older experimental work by, for example, (Smith (1962)) seems to confirm these theoretical predictions. These results are relevant for the analysis of such trading institutions as stock exchange, as well as internet exchanges which use integrated trading systems and match buyers and sellers through a single book of orders¹.

However, a host of new trading institutions, in particular internet auction markets, such as EBay and Ubid, and business-to-business exchanges that have emerged recently² do not fit into the standard models of centralized multilateral exchange, because they are not governed by any centralized pricing mechanism. Although these trading institutions provide marketplaces that serve as focal location for trade and set out basic rules of trading, unlike in double auction markets and in auctions with a single seller, they do not process demand and supply information, and leave the responsibility for coordinating trade to the traders themselves. In particular, it is up to the sellers to choose their own reserve prices, whereas buyers decide where to submit bids. Experienced traders apparently buy and sell in many locations.

These trading institutions are made possible by the progress of modern computer technology which allows individual traders to collect dispersed information about prices and availability and respond to it promptly.³ The volume of trade that they intermediate is growing at a fast rate. 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)).

One of the central issues that these institutions bring to the forefront is whether efficiency can be achieved under decentralization. This issue has been previously addressed in other contexts. For example, in the context of random matching Rubinstein and Wolinsky (1990) demonstrate that bargaining may or may not produce a competitive outcomes when trade frictions are small. In this paper, we extend this line of inquiry to internet auctions and exchanges. We model them as a market with very low trading frictions where a sufficiently large number of buyers and sellers interact. The sellers independently run ascending second-price auctions similar to those that take place on EBay. Buyers submit bids at these

¹For example, the NTE truck exchange (<http://www.nte.net>) allows truckers with excess carrying capacity to match with customers who need to ship small and unusual packages. 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 will attract too few buyers and sellers. Instead, supply and demand information for related routes and times is aggregated and used to generate market prices.

²The examples of large-scale business-to-business internet exchanges include CheMatch and Chem Connect (chemical industry), Covisint (auto parts), MetalSpectrum (aluminum, stainless steel, copper, iron, and other metals).

³Arguably, the widespread use of personal computers enhances decentralized decision-making as much as it does allow to decentralize computation. As a further step towards decentralization, one could imagine that the robot guided bidding, such as the one used on EBay, could be carried out on a host of small private servers (owned by sellers) which would simply enforce the trading rules of a larger institution.

auctions. If their bids are not successful, the buyers can adjust them and move between sellers costlessly.

Thus, trade is completely decentralized, and the trading rules are very similar to the ones that are used in well-known internet auction markets, such as EBay and other business-to-business marketplaces.⁴

Generally, the outcome of decentralized trading will be efficient if two conditions are satisfied. First, the rule which is used to compute prices and efficient trades can be broken up into bits which can be solved by individual traders. Second, the traders have an incentive to follow this pricing rule. Accordingly, the key part of this paper lies in the analysis of the bidding rule - the strategy that buyers use to select among sellers' auctions and choose their bids, and of the reserve price strategy used by the sellers.

The bidding rule that we design is a simple function of the publicly observable market data. It requires a buyer to bid at an auction with the lowest current price and raise his bid as slow as possible (as long as a buyer's bid is below his valuation). Otherwise, it is independent of the buyers' beliefs about costs or valuations and even the number of other buyers who participate in the process. This bidding behavior will tend to equalize prices across different auctions. Some price dispersion may still be possible as a result of miscoordination. However, as we demonstrate, price dispersion is avoided if the buyers follow a simple priority rule when choosing among the auctions with the lowest current price. They should first of all choose among the auctions where no bids have been received or where there has been more activity (in the sense which will be made precise below) after the last change of the winning bidder.

We demonstrate that the specified bidding rule constitutes a Bayesian Nash (continuation) equilibrium for buyers given the sellers' reserve prices. When buyers follow this bidding rule, all trades will occur at a uniform price equal to the highest valuation among buyers who fail to win an auction or to the highest reserve price among successful sellers. This outcome is the same as in "seller's offer" double auction where all bids and asks are processed in the centralized way⁵. If sellers' reserve prices happened to be equal to their true valuations, then the uniform trading price would be equal to the lowest market clearing price.

In the second part of the paper, we show that it is a Bayesian equilibrium for sellers to set their reserve prices equal to their true costs if the number of buyers and sellers in the market is large (but finite), and costs and valuations are distributed independently. Then efficient trades occur at a uniform price. In this sense the decentralized market supports 'competitive outcome'⁶.

Other results demonstrating that efficient outcomes can be implemented via an auction mechanism (see Gul and Stacchetti (2000), Ausubel (2000), Roth and Sotomayor

⁴In the case of Ebay, the bids are actually processed on EBay servers. However, our model also describes markets where the sellers conduct the same auctions on their own servers. In fact, the boundaries of any online market go far beyond the confines of a single server.

⁵For the definition of the "seller's offer" double auction see Satterthwaite and Williams (1989b).

⁶We get stronger results than Rustichini, Satterthwaite, and Williams (1994) concerning the existence and efficiency of equilibria because we assume that there is a finite set of feasible valuations rather than a continuum. This assumption also allows us to get around the impossibility result of Myerson and Satterthwaite (1983).

(1990)) have been obtained in environments that are different from the one that we study. The interest in our procedure stems from the fact that the price adjustments are done by traders themselves without any need for centralized intervention. Also, the strategies that the traders use are simple, and very little information is required to carry them out. The traders do not need to perform sophisticated calculations that otherwise might have to be performed by the mechanism designer intermediating the trades, or by buyers and sellers attempting to foresee future trading prices in a rational expectations model. Buyers need to know only the current set of prices (second-highest bids), and whether or not the current price at a given auction has changed recently (more precisely, after the last change of the winning bidder).

The requirement that individual sellers run second-price auctions and the observability conditions can be interpreted as rules imposed by the large institution, or the exchange, in the context of which individual sellers offer their mechanisms. Although we do not offer a formal analysis of the optimality of these rules, we provide an argument suggesting that these rules are more efficient than other plausible alternatives, such as first-price auctions and mechanisms where the winning bids are observable. In particular, the current rules motivate the buyers to bid as high as their valuations, and make preemptive bidding (i.e. submitting very high bids early in the game) unprofitable.

Apart from their theoretical value, the results of this paper have potentially important practical implications since a decentralized market is less prone to technological failures⁷.

If buyers follow the strategies that we describe, predictable patterns of bidding and prices will emerge. In particular, the winning bid in our model is not necessarily the last bid. However, because the bidders push up the standing bids slowly, the winning bid will normally be submitted towards the end of an auction. It would be rare in our mechanism for the winning bid to be submitted early (in the sense that there are many bids submitted after it). There is strong evidence that EBay auctions are won by late bidders (Zheng (2001) or Roth and Ockenfels (2000)).

It is readily apparent in our model why buyers do not want to follow EBay's advice and bid their true valuations right away, leaving the robots to compute the prices later. Buyers with sufficiently high valuations will win the auctions that they participate in, whether they bid early or late. However, a buyer who bids his true valuations early may accidentally be trapped into paying a higher price than necessary, when another buyer with a relatively high valuation bids against him. Raising the bid as slowly as possible (by submitting the minimum acceptable bid) avoids this.

Another feature of the equilibrium is that winning bidders will tend to do more bidding than losing bidders (not necessarily all in the same auction). There is some preliminary evidence which, consistently with our prediction, indicates that bidding activity and success are positively correlated. Our model also predicts strong negative correlation between seller's reservation price and bidding activity. Sellers with high reservation prices should get few bids. However, at least to the extent that the commodities being sold are substitutes, our model also predicts that the units with high and lower reservation price

⁷The Toronto Stock Exchange for example, has suffered continuing technical problems that have shut down its computerized trading system repeatedly and resulted in millions of dollars in trading losses.

should end up trading at the same price - so the correlation between trading price and the seller's reserve price should be weak.

There are at least two branches in the literature that are related to our paper. First, we contribute to the literature on competing auctions including McAfee (1993), Peters (1997), Peters and Severinov (1997), Hernando-Veciana (2000). The flavor of the results in that literature is similar to our results - when there are many competing auctions, *equilibrium* mechanisms are independent of the characteristics of demand, and second-price auctions with reserve prices set equal to costs are optimal *mechanisms* for sellers independently of their beliefs. We also find that it is optimal for sellers to set their reserve prices equal to true costs (though sellers are only free to choose a reserve price in our model, they cannot use mechanisms other than auctions). The difference is that in the previous literature the cost of moving between sellers is implicitly very high in the sense that buyers are able to bid at only one seller. As a consequence, the equilibrium outcomes in these models are not ex post efficient. Coordination problems are large enough that there are unrealized gains from trade. In our model bidders can, at zero cost, submit as many bids at as many sellers as they like. In effect, this costless communication allows buyers to get around the coordination problems that arise in the earlier models.

We have already mentioned other results on implementation of efficient outcomes via auction mechanisms, especially Gul and Stacchetti (2000), Ausubel (2000) who allow for heterogeneous commodities in single seller environments. Milgrom (2000) studies centralized ascending auction with heterogeneous objects and discusses the evidence from recent spectrum auctions, and the effects of various bidding rules on buyers' behavior. Roth and Sotomayor (1990) describe an auction-like mechanism which generates competitive trade in an environment similar to the one that we consider. The incentive properties of this mechanism were studied earlier by Demange (1982) and Leonard (1983). Their mechanism is not decentralized in the sense that ours is, since it involves an auctioneer who collects information and adjusts prices⁸.

Our bidding procedure is also close in spirit to the 'deferred acceptance' algorithms of Gale and Shapley (1962) and Kelso and Crawford (1982) (the latter study the case with transferable utility). These algorithms are also decentralized to a certain extent, since proposal acceptance decisions can be left up to individual traders. It is known that the deferred acceptance algorithms are not strategy-proof when utility is non-transferable. Less is known about the incentive properties of these algorithms with transferable utility.

The rest of the paper is organized as follows. In the next section we describe our model. In section 3 we provide a description of the buyers' strategies, and show that these strategies constitute a Bayesian Nash equilibrium in the continuation game after the sellers post their reserve prices. In section 4 we analyze the sellers' optimal reserve prices. The equilibrium outcome is the same as in "seller's offer" double auction. In fact, our argument

⁸Roughly each buyer tells the auctioneer which sellers he would be willing to purchase from at the current set of prices. Given this information, the auctioneer checks whether it is feasible to match buyers to sellers. If not, then prices are increased at all the sellers where there is an excess demand. Buyers can send messages to each of the sellers in their 'demand sets' informing them that they are willing to trade, but sellers who have an excess demand cannot tell whether a feasible match is possible or whether they should raise price without a signal from the auctioneer.

shows that "seller's offer" double auction has an efficient equilibrium when the set of costs and valuations is finite. This could be viewed as an independent contribution of the paper.

2 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 $\mathcal{D} \equiv \{\underline{d}, \underline{d} + d, \underline{d} + 2d, \dots, \bar{d}\}$ that has a step size $d > 0$. Let $F(\cdot)$ and $G(\cdot)$ be the probability distributions from which a buyer's valuation and a seller's costs 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 $v - p$. 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 seller with cost c who sells at price p gets 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} .⁹ The grid size d could be thought of as the minimum bid increment. When a 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. We also assume that, as in most on-line auctions, the buyers also know the identity of the high (winning) bidder at each seller at all times¹⁰. The standing bids and the identities of the winning bidders are the only two observables in our model.

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

⁹This 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.

¹⁰As will be shown later, this assumption, or more precisely, the observability of the change in the identity of the high bidder is important for the uniform price result. When such changes are not observable, price dispersion may occur.

last bidder followed by the first bidder and so on. Bidding continues until all buyers pass. Then 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.

The goal of this paper is not so much to model the details of a particular auction mechanism, as it is to establish the existence of a relatively simple decentralized process that generates an ex-post efficient outcome. It seems reasonable to conjecture that the same result holds under a large class of auction rules. However, the bidding game that buyers play is complex enough that verifying the robustness of our results to small changes in the auction rules is quite difficult.


Despite the second price nature of the auction mechanism, the presence of multiple auctions implies that it is not a dominant strategy and not even a Nash equilibrium strategy for buyers to bid their true valuations when they start bidding (despite the advice offered on the EBay website). To see this, consider a market with two buyers and two sellers. Let b_1 (b_2) be buyer 1's (2's) valuation, and suppose that sellers 1 and 2 announce reserve prices s_1 and s_2 respectively. Let $b_1 > s_2 > s_1$. Suppose that buyer 1 enters first, and that buyer 2 is expected to bid his valuation at the seller with the lower standing bid. Visually, consider data in Figure 1, but ignore point b_3 and allow b_2 to take all possible values.

We will show that buyer 1 can earn a higher expected payoff if, instead of bidding her true valuation b_1 at the start of the bidding game, he submits a bid equal to s_2 at seller 1. If buyer 1 starts by bidding his valuation with seller 2, then buyer 2 will 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 . Yet, buyer 1 can strictly increase her expected payoff by initially bidding s_2 at seller 2, because there is a positive probability that buyer 2's valuation is below s_2 .

Next, suppose that buyer 1 bids his valuation with seller 1. When buyer 2 enters, the standing bid with seller 1 is still equal to s_1 , because seller 1 has yet to receive a second bid. So, according to buyer 2's strategy, she will bid her valuation at seller 1. If buyer 2's valuation is s_2 or less, then buyer 1 will trade at price equal to buyer 2's valuation. The same outcome will occur if initially buyer 1 bids s_2 , instead of b_1 , at seller 2. If buyer 2's valuation is above b_1 , then buyer 1 will be displaced as high bidder at seller 1 and will trade at price s_2 with seller 2, no matter whether she bids b_1 or s_2 initially.

Finally, suppose that buyer 2's valuation lies in the interval $[s_2, b_1]$ (which occurs with a strictly positive probability). If buyer 1 bids b_1 at the start, she will trade with seller 1 at a price equal to buyer 2's valuation. However, if buyer 1 starts by bidding s_2 with seller 1, then she will be displaced as high bidder at seller 1 and will end up trading at price s_2 with seller 2. Thus, bidding s_2 ensures that buyer 1 never has to pay a price above s_2 and improves her expected payoff.

Figure 1:





3 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 bid, have an option to bid again elsewhere. Unfortunately, this option is not sufficient to guarantee that, conditional on sellers' 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 $\mathbf{b} = \{b_1, \dots, b_m\}$ be the vector of buyers' valuations, and let $\mathbf{s} = \{s_1, \dots, s_n\}$ be the vector of sellers' reserve prices. Let $\mathbf{v} = \{v_1, v_2, \dots, v_{m+n}\}$ be vector with entry v_i equal to i -th lowest value (from the bottom) among buyers' valuations and sellers' reserve prices. If sellers' reserve prices are 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 less than or equal to 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.

We now define the symmetric strategy σ^* by specifying how each buyer should bid when his turn comes.

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 with lowest standing bid who has not yet received a bid, or where the standing bid has changed since the last change of the winning bidder. If all sellers with lowest standing bid have received bids, and at all of these sellers the standing bid has not changed since the current high bidder has assumed her position, then the buyer should bid with each of them with equal probability.*

If all active buyers follow strategy σ^* , then irrespective of their valuations and beliefs, they would focus their bids on the group of sellers (which could consist of only one seller) with the lowest standing bids. The active buyers continue to bid up the standing bid with this group of sellers until their standing bids reach the level of the standing bids at the next group of sellers (which in equilibrium would be equal to the reserve price of these sellers). Then bidding continues with sellers from both groups until the lowest standing bid reaches the level of 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.

The only part of strategy σ^* that requires a more detailed explanation is rule (c). The sellers with the same standing bid can have different winning bids. Although winning bids are unobservable, rule (c) allows the buyer to identify which of the sellers with the

lowest standing bid have lower winning bids. Bidding only at such sellers is optimal for the buyer, because it increases his chances of trading at a lower price. At the same time, such bidding ensures that standing bids rise uniformly and eliminates the possibility of price dispersion.

To understand how rule (c) works, note that on the path generated when all buyers use strategy σ^* a new winning bid will be strictly above the standing bid (by one grid point) and will displace the previous winning bidder without changing the standing bid. If another buyer later raises the standing bid without displacing the high bidder, the standing bid becomes equal to the winning bid. However, if standing bid does not change after a change in the identity of the winning bidder, then the winning bid remains above the standing bid.

It is easy to see that the following alternative rule achieves the same result as rule (c). A buyer should choose, first of all, among the sellers with lowest standing bids who have not yet received any bids or where the last bid has caused a change in the standing bid¹¹. An intuitive explanation to rule (c) and its alternative can be provided in the following way. If a winning bidder has been "attacked" previously by other buyers whose bids have caused an increase in the standing bid, then it is likely that the standing bid is close or equal to the winning bid, and therefore a buyer's chances of overtaking the winning bidder are higher.

Before proceeding further, 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 indexes, so buyer 3 with valuation b_3 enters first, and buyer 1 enters last. 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 also bid s_2 with seller 1, since seller 1's standing bid is still the lowest. This bid by buyer 2 will not be a winning bid, 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 bid s_2 , 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 standing bid with seller 1 has changed since the winning bidder changed, 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 will become high bidder but his bid will not affect either of the standing bids. Note that a winning bid will not change a seller's standing bid on the equilibrium path.

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, but does not affect either of the standing bids which remain equal to s_2 . When

¹¹Buyers are able to detect when a seller receives a bid, because each new bid changes either the winning bidder or the standing bid.

buyer 1 enters he finds that the standing bid has not changed since winning bidder changed at either seller, so he will bid b_3 at one seller, and then at the other in random order. Neither bid by seller 1 will be successful, but 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 standing bids have increased since the winning bidder has changed at both sellers). If buyer 1 bids b_2 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 buyer 1 bids b_2 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 combinations 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 1 *The outcome in which all buyers use the strategy σ^* is a perfect Bayesian equilibrium. Given the 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 price v_m .*

Proof: see the appendix.

It is notable that strategy σ^* contains the same set of rules on the equilibrium path 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^* can at best raise the trading price. At the same time, a buyer who bids as if his valuation is lower than what it really is, 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 optimal in a standard static second price auction. But the analogy is only approximate, because we have to deal with a dynamic bidding process and presence of multiple auctions¹². To compute the payoff that a buyer gets by following σ^* from the

¹²As shown above, bidding one's valuation is not an optimal strategy in the dynamic bidding game.

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^* .

According to theorem 1, in the equilibrium where buyers use σ^* all trades will be completed at uniform price equal to v_m , m -th lowest value (from the bottom) in the array of buyers' valuations and sellers' costs. The uniform price result is due to rule (c) in σ^* . By following this rule, a buyer is able to focus his bidding on those sellers with lowest standing bid who also have lower winning bids. In the end, such bidding also eliminates the possibility of price dispersion.

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 and $k - 1$ may trade in equilibrium.

To better understand whether buyers and sellers whose valuations and costs are equal to v_m trade, let us divide the sets of buyers and sellers into three groups. For given array \mathbf{v} , let $M_1/M_2/M_3$ be the set of buyers whose valuations are, respectively, *strictly lower than* v_m /*exactly equal to* v_m /*strictly higher than* v_m . Similarly, let $N_1/N_2/N_3$ be the sets of sellers who set their reserve prices below/equal to/ above v_m respectively. Let m_i (n_i) be the number of buyers (sellers) in the set M_i (N_i). Theorem 1 says that buyers in M_3 will surely trade, and that buyers in M_1 and sellers from N_3 will not trade.

Corollary 1 *Let v_m be the m -th lowest element in the array \mathbf{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*

$$\max[0, m_3 - n_1] \quad \text{and} \quad \min[n_2, m_3 - \min\{0, n_1 - m_2\}]$$

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

$$\max[0, n_1 - m_3] \quad \text{and} \quad \min[m_2, n_1 - \min\{0, m_3 - n_2\}]$$

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).¹³ But no matter

¹³By 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 1. To obtain the

how many sellers from N_2 and buyers from M_2 trade, this has no effect on the traders from other groups 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.

The central implications of theorem 1 and corollary 1 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. Thus, 'apparent gains from trade' are maximized.

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.¹⁴ For example, when there is only one seller, the trade will occur at the reserve price if there is only one buyer whose valuation is above the reserve, and at the second highest valuation otherwise. This is the same outcome as in a second-price auction. Playing strategy σ^* remains an equilibrium when there is a large number of sellers as well. Buyers' behavior is efficient because buyers affect the trading price only when they forego attractive trading opportunities.

Of course, efficiency in the buyers' game does not guarantee that the equilibrium outcome of the mechanism will be efficient, because sellers may set reserve prices that are different from their true costs. We turn to examination of sellers' behavior in the next section.

We conclude this section by pointing out that the efficiency in the bidding game can be attributed to the rules of our mechanism which require individual sellers to run second-price auctions and the observability conditions according to which a buyer observes only standing bids and the identities of the winning bidders. In our interpretation, these rules are imposed by a large institution (internet exchange) in the context of which individual sellers offer their goods. Thus, it is important to understand what the consequences might be if this institution used a modified set of rules.

The argument that we provide is informal and partial,¹⁵ yet, despite these shortcomings it is quite useful for our discussion. It suggests that the rules of our mechanism possess certain desirable optimality properties. To show this, we will compare the equilibrium outcome of our mechanism to plausible equilibria in plausible alternative mechanisms.

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 who already have buyers as winning bidders (as opposed to sellers who have not yet received any bids). 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.

¹⁴This 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 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.

¹⁵We focus on specific equilibria and do not consider a sufficiently general class of alternative mechanisms.

First, suppose that sellers had to use first-price auctions. It is easy to see that buyers will never bid their true valuations because in this case their payoffs would be zero. Therefore, inefficiency will arise in the bidding game. When the number of traders is large, this inefficiency cannot be mitigated by sellers' decision to set lower reserve prices in the first-price auctions than in the second-price auctions. This is so, because in the next section we demonstrate that sellers will set reserve prices equal to their true costs in our mechanism when the number of traders is sufficiently large. Note that buyers are willing to bid up to their valuations in our mechanism because they still hope to trade at a lower price.

Second, the non-observability of the winning bids ensures that preemptive bidding does not occur in our mechanism. In contrast, suppose that winning bids are observable, but the second-price auction format is intact. Then there is an equilibrium in which the first buyer in order of entry submits the highest possible bid at the seller with the lowest reserve price, the second buyer in order of entry submits the highest possible bid at the seller with the second lowest reserve price, and so on. No other buyer bids at a seller who has received the highest possible bid because no other buyer will be able to trade with such seller. Thus, buyers who are among the first to enter are able to 'clinch' the goods by such preemptive bidding and trade at the reserve prices. Since these buyers can have lower valuations, inefficiency arises. Clearly, in this equilibrium sellers will be more inclined to post reserve prices above costs than in the equilibrium that we study.

4 Optimal Reserve Prices in a Large Market

There is a notable connection between our result and equilibrium in the double auction market. Consider a "seller's offer" double auction where uniform trading price is set equal to m -th lowest value (from the bottom) among buyers' bids and sellers' asks. Theorem 2.1 in Satterthwaite and Williams (1989a) can be extended in a straightforward way to show that buyers would bid their true valuations in this double auction. Bidding his valuations is optimal for a buyer because his bid can only affect trading price if the buyer fails to trade. The usual second price logic then applies. Thus, the trading price in both 'seller's offer' double auction and in our decentralized market is equal to m -th lowest value in the array of buyers' valuations and sellers ask/reserve prices.

Further, because in our decentralized market and in a double auction all sellers set their reserve prices simultaneously, a seller's strategic situation is the same in both mechanisms. For fixed array of valuations and given strategies of other sellers, a seller gets the same expected payoff from her action in both mechanism. Therefore, the equilibrium outcome of decentralized mechanism studied here is identical to an equilibrium outcome in "seller's offer" double auction.

Double auctions are known to be incentive efficient (i.e. second-best) when the number of traders is sufficiently large (see Wilson (1985)). In this section, we establish an even stronger result due to the fact that we consider a finite grid of valuations: ex post efficient outcome is achieved when there are many buyers and sellers.

Turning to the sellers' decision, our goal is to show that in a large market it is an equilibrium for sellers to set reserve prices equal to their true cost. This result is similar to that of Satterthwaite and Williams (1989a) and Rustichini, Satterthwaite, and Williams

(1994) who analyze the rate of convergence of sellers' ask prices and buyers' bids to their true costs and valuations in a double auction 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.

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 $\bar{f} > 0$.

It may not be optimal for a seller to set a reserve price equal to her true valuation, because her reserve price may in some cases affect the trading price. For example, consider the situation depicted in Figure 1 without buyer 3 (the one with the lowest valuation) and let buyer 2's valuation b_2 still be undetermined. 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 . Since b_2 is random, seller 2 could raise the trading price with a strictly positive probability by raising her reserve price above s_2 . The cost of doing this to seller 2 is that she would fail to trade if buyer 2's valuation happens to be between s_2 and her new higher reserve price. This tradeoff is similar to the one which traders face in the standard double auction.

Let us consider a sequence of auction 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 2 *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, seller j with cost c obtains a higher expected payoff by setting reserve price equal to $p-d$ rather than p where $p > c$. When the trading price happens to be above p or below $p-d$, it does not matter whether the seller posts price p or $p-d$. Thus, we need to focus on the situations where seller j 's choice between p and $p-d$ affects the trading price, or where the trading price is equal to $p-d$ or p irrespective of seller j 's choice. The latter situation occurs if some other sellers post the reserve price equal to the trading price or some buyers' valuation are equal to it. 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. If several sellers post reserve prices equal to the trading price, several of them may fail to trade because of the competition from the others.

Precisely, seller j will get different payoffs from setting her 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 whether the seller sets her reserve price equal to $p - d$ or p , and seller j fails to trade after setting reserve price p .
3. 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 ¹⁶.
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 (3) and show that the effect of (2) alone dominates (4). The effects of (1) and (3) are ignored because when $p - d = c$, the net surplus to seller j from trading at price $p - d$ is zero. Comparing (2) 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. Competition from these sellers implies that a seller posting reserve price equal to trading price will often fail 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.

5 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 intend 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. At the same time, the equilibrium behavior that we identify has a number of advantages. It is simple, requires very little computation on the part of traders, and is invariant to the distributions from which costs and valuations are drawn. It also implements the efficient allocation.

¹⁶It is easy to show that seller j will always trade after posting reserve price $p - d$ in this case

These properties make it reasonable to 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 aspects and rules 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.

As mentioned above, our model gives a very simple explanation for the fact that buyers do not bid their true valuations. Patient revision of bids provides buyers with the signals that they need in a decentralized market to coordinate their bidding behavior. Deviating from this pattern will often force a buyer to trade at a price that is higher than necessary.

The empirical implications of our theory are discussed in the introduction. An extension of the logic of our model suggests an explanation for observed flurry of active bidding close to the end dates in the EBay auctions. In our model, all the available trading opportunities are visible at the beginning of the bidding process. On EBay, these opportunities arise at random times. One of the opportunity costs of submitting a bid on EBay arises from the possibility that a new seller will enter and post a lower reserve price after the buyer submits his bid. The buyer may then end up trading at a higher price than he needs to. Effective coordination of bidding among buyers then demands that they refrain from starting the bidding as long as possible. Since auctions at EBay have fixed end dates, at some point the probability that new sellers will enter before the current auction ends becomes small, and buyers will start bidding. Our theory suggests that this late bidding on an auction that is ending will induce a flurry of bidding at other auctions for similar goods.

6 Appendix

Proof of Theorem 1:

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 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 (from the bottom) 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_Γ} .

Lemma 1 *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_Γ} .*


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 (from the bottom). 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_Γ} (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 yet. The standing bids remain s_1 at

Figure 3:



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 2 *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_Γ} .*

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_Γ} . By the definition of \tilde{v}_{m_Γ} , there must be at least k sellers whose standing bids in state Γ are \tilde{v}_{m_Γ} 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 any 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_Γ} 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 standing bid 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 standing bid at j has not changed since i became the winning bidder, buyer i' will bid at j only if s^j is the lowest standing bid, and at any of the sellers with standing bid s_j , including sellers whose standing bids are below \tilde{v}_{m_Γ} in state Γ , standing bid has not changed since winning bidder has changed.

This implies that all k sellers who in state Γ have standing bids below \tilde{v}_{m_Γ} 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_Γ} 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 $b^i > \tilde{v}_{m_\Gamma}$ 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 $s^{j'} = \tilde{v}_{m_\Gamma}$ and $p = b^i = \tilde{v}_{m_\Gamma} + d$. Consistency of i 's position with σ^* then implies that standing bid has not changed at j since i has become a winning bidder at standing bid $s^{j'}$.

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

Lemma 3 *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_Γ} . 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 at a seller where the standing bid 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 becomes 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 4 *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 high 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 3, 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 3- 4 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, positions of all buyers 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 5 *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 deviating from σ^* provided that the other buyers follow σ^* .*

Proof. By lemmas 3 and 2, 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 3 along any such path can never be lower than v_{m_Γ} . ■

To complete the description of the equilibrium, we specify buyers' beliefs on and off the equilibrium path. 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 the standing bids at sellers where she is a high bidder, (ii) the game is in a regular state. Further, if standing bid at a given seller has changed after

the last change in the identity of the winning bidder, then all buyers except the winning bidder at this seller believe that with probability 1 the winning bid is equal to the standing bid. If standing bid has not changed since the last change of the winning bidder, then all buyers except the winning bidder at this seller believe that the winning bid is one step (d) higher than the winning bid.

Thus, buyers hold beliefs that are consistent with all buyers playing σ^* to the extent that this is possible. Certain deviations from σ^* , for example, a buyer bidding above his valuation, are not observable and do not affect other buyers' beliefs. Deviations that are detected are of two kinds. First of all, a buyer may become high bidder with more than one seller. Beliefs in this case are described by (i). Further, if standing bid at a seller changes more than once without a change in the identity of the winning bidder, then the winning bidder must have deviated from σ^* . The buyers in this case maintain the beliefs that the deviation is the minimum possible one consistent with the path of the game, i.e. the winning bid is equal to the standing bid.

The proof that σ^* together with the specified beliefs constitute a Bayesian equilibrium in the bidding game given sellers' announced reserve prices then follows immediately from Lemma 5 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 1:

Suppose that seller j posts reservation price s s.t. $s < v_m$. According to theorem 1, 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 randomization 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 2:

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. The latter are equal to the true costs for all sellers except z by assumption of the theorem. 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 . Seller z 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 $p < v_m$. By corollary 1, every seller who posts a price below v_m trades.

Claim 2. *Suppose that if seller z posts reserve price p , then the trading price is p_T s.t. $p > p_T$. Then the trading price will also be p_T , if seller z posts reserve price $p'' > p_T$. Seller z 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 $p > v_m$. By corollary 1, any seller who posts a price above v_m does not trade.

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 2 to hold.

Lemma 6 *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:

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

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, if there are other sellers who post this reserve price or 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:

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

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

■

Lemma 7 *Inequality (1) holds if:*

$$\begin{aligned} &E\left(\frac{1 + n'_2 + n_1 + m_1 - m}{n'_2 + 1} \mid m_1 + n_1 < m - 1 \leq m_1 + n_1 + m_2 + n'_2\right) \times \\ &\times P(m_1 + n_1 < m - 1 \leq m_1 + n_1 + m_2 + n'_2) \\ &\geq E\left(\frac{1}{n'_2 + 1} \mid m_1 + n_1 = m - 1\right)P(m_1 + n_1 = m - 1) \end{aligned} \quad (3)$$

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 m_3 buyers whose valuations are above p . By corollary 1, 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

$$P^t \equiv \min\left\{1, \frac{m_3 + \min\{0, m_2 - n_1\}}{n'_2 + 1}\right\}$$

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

$$P^f \equiv \max\left\{0, \frac{n'_2 + 1 - m_3 - \min\{0, m_2 - n_1\}}{n'_2 + 1}\right\}$$

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

$$A_1(\alpha) \equiv \sum_{\hat{m}_1 = [\alpha m] + 1}^{m-1} E\left(\frac{1}{n'_2 + 1} | n_1 = m - 1 - \hat{m}_1\right) P(m_1 = \hat{m}_1) P(n_1 = m - 1 - \hat{m}_1)$$

$$A_2(\alpha) \equiv \sum_{\hat{m}_1 = \max\{0, m - n\}}^{[\alpha m]} E\left(\frac{1}{n'_2 + 1} | n_1 = m - 1 - \hat{m}_1\right) P(m_1 = \hat{m}_1) P(n_1 = m - 1 - \hat{m}_1)$$

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:

$$\begin{aligned}
B_1 &\equiv \sum_{\hat{m}_1=\max\{0,m-n\}}^{m-2} E\left(\frac{1+n'_2+n_1+m_1-m}{n'_2+1}\middle| m_1=\hat{m}_1+1, n_1=m-\hat{m}_1-3\right) \times \\
&\quad \times P(m_1=\hat{m}_1+1)P(n_1=m-\hat{m}_1-3) \\
B_2 &\equiv \sum_{\hat{m}_1=\max\{2,m-n+4\}}^{m-1} E\left(\frac{1+n'_2+n_1+m_1-m}{n'_2+1}\middle| m_1=\hat{m}_1-2, n_1=m-\hat{m}_1-1, n'_2 \geq 4\right) \times \\
&\quad \times P(m_1=\hat{m}_1-2)P(n_1=m-\hat{m}_1-1, n'_2 \geq 4)
\end{aligned}$$

To complete the proof of the theorem, we will establish the following lemma:

Lemma 8 $\exists \alpha \in (0, 1)$ s.t. if and m (and hence n) is large enough then $A_1(\alpha) < B_1$ and $A_2(\alpha) < B_2$.

Proof. First, note the following expressions. Let $a+b+c=n-1$ and $r+s+t=m$. Then:

$$P(n_1=a, n'_2=b, n_3=c) = \frac{(n-1)!}{a!b!c!} G(p-1)^a g(p)^b (1-G(p))^c \quad (4)$$

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

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

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

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

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

Step 2. Suppose that $n_1 = m - \hat{m} - 1$. 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, $\hat{m}_1 \geq \alpha m$ and therefore $n_2 + n_3 \geq \frac{m}{2k}$. Since $g(p) \geq \bar{g} > 0$, by the law of large numbers $\exists m''$ s.t. if $m > m''$ then $P(n'_2 \geq 4 | n'_2 + n_3 \geq \frac{m}{2k}) > 5/6$.

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,

$$\begin{aligned}
&E\left(\frac{1}{n'_2+1}\middle| n_1=m-\hat{m}_1-1\right)P(n_1=m-\hat{m}_1-1) = \\
&E\left(\frac{1}{n'_2+1}\middle| n'_2 < 4, n_1=m-\hat{m}_1-1\right)P(n'_2 < 4, n_1=m-\hat{m}_1-1) + \\
&E\left(\frac{1}{n'_2+1}\middle| n'_2 \geq 4, n_1=m-\hat{m}_1-1\right)P(n'_2 \geq 4, n_1=m-\hat{m}_1-1) \\
&< (1 \times 1/6 + 1/5 \times 5/6) P(n_1=m-\hat{m}_1-1) = 1/3 P(n_1=m-\hat{m}_1-1) \quad (7)
\end{aligned}$$

On the other hand, we have:

$$\begin{aligned} & E \left(\frac{1 + n'_2 + n_1 + m_1 - m}{n'_2 + 1} \middle| m_1 = \hat{m}_1 - 2, n_1 = m - \hat{m}_1 - 1, n'_2 \geq 4 \right) P(n_1 = m - \hat{m}_1 - 1, n'_2 \geq 4) \\ & \geq \frac{2}{5} P(n_1 = m - \hat{m}_1 - 1, n'_2 \geq 4) > \frac{1}{3} P(n_1 = m - \hat{m}_1 - 1) \end{aligned} \quad (8)$$

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 n is sufficiently large).

$$\begin{aligned} & E \left(\frac{n_2 + 1 + n_1 + m_1 - m}{n_2 + 1} \middle| n_1 = \hat{n}_1 - 2, m_1 = \hat{m}_1 + 1, n_2 = \hat{n}_2 + 2 \right) \times \\ & \times P(n_1 = \hat{n}_1 - 2, m_1 = \hat{m}_1 + 1, n_2 = \hat{n}_2 + 2) \\ & > \frac{1}{n_2 + 1} P(n_1 = \hat{n}_1, m_1 = \hat{m}_1, n_2 = \hat{n}_2) \end{aligned} \quad (9)$$

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

$$\begin{aligned} & E \left(\frac{n_2 + 1 + n_1 + m_1 - m}{n_2 + 1} \middle| n_1 = \hat{n}_1 - 2, m_1 = \hat{m}_1 + 1, n_2 = \hat{n}_2 + 2 \right) \times \\ & P(n_1 = \hat{n}_1 - 2, m_1 = \hat{m}_1 + 1, n_2 = \hat{n}_2 + 2) = \frac{\hat{n}_2 + 1}{\hat{n}_2 + 3} P(n_1 = \hat{n}_1 - 2, m_1 = \hat{m}_1 + 1, n_2 = \hat{n}_2 + 2) \\ & = \frac{\hat{n}_2 + 1}{\hat{n}_2 + 3} P(n_1 = \hat{n}_1, m_1 = \hat{m}_1, n_2 = \hat{n}_2) \frac{g(p)^2}{G(p-1)^2} \frac{(\hat{n}_1 - 1)\hat{n}_1}{(\hat{n}_2 + 2)(\hat{n}_2 + 1)} \frac{F(p-1)}{1 - F(p-1)} \frac{m - \hat{m}_1}{\hat{m}_1 + 1} \\ & = \frac{1}{\hat{n}_2 + 1} P(n_1 = \hat{n}_1, m_1 = \hat{m}_1, n_2 = \hat{n}_2) \frac{g(p)^2}{G(p-1)^2} \frac{F(p-1)}{1 - F(p-1)} \frac{(\hat{n}_2 + 1)(\hat{n}_1 - 1)\hat{n}_1}{(\hat{n}_2 + 2)(\hat{n}_2 + 3)} \frac{m - \hat{m}_1}{\hat{m}_1 + 1} \end{aligned}$$

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 $\hat{n}_2 \leq (\frac{1}{k} + \alpha - 1)m$. Therefore,

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

The right-hand side of the above inequality is 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. (2000): "An Efficient Dynamic Auction for Heterogeneous Commodities," University of Maryland.

- DEMANGE, G. (1982): "Strategyproofness in the Assignment Market Game," Laboratoire d'Econometrie de l'Ecole Polytechnique, Paris.
- GALE, D., AND L. SHAPLEY (1962): "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, 69, 9–15.
- GUL, F., AND E. STACCHETTI (2000): "The English Auction with Differentiated Commodities," *Journal of Economic Theory*, 92(1), 95–124.
- HERINGS, J. J. (1997): "A Globally and Universally Stable Price Adjustment Process," *Journal of Mathematical Economics*, 27, 163–193.
- HERNANDO-VECIANA, A. (2000): "Competition Among Auctioneers," Discussion paper, Universidad de Alicante.
- KELSO, A. S., AND V. P. CRAWFORD (1982): "Job Matching, Coalition Formation and Gross Substitutes," *Econometrica*, 50, 1483–1504.
- LEONARD, H. B. (1983): "Elicitation of Honest Preferences for the Assignment of Individuals to Positions," *Journal of Political Economy*, 91, 461–79.
- MCAFEE, P. (1993): "Mechanism Design by Competing Sellers," *Econometrica*, 61(6), 1281–1312.
- MILGROM, P. (2000): "Putting Auction Theory to Work: The Simultaneous Ascending Auction," *Journal of Political Economy*, 108(2), 245–272.
- MORNEAU, J. (2001): "E-Marketplaces Start to Take Off," *TechWeb News*, <http://www.techweb.com>.
- MYERSON, R., AND M. SATTERTHWAIT (1983): "Efficient Mechanisms for Bilateral Trading," pp. 265–281.
- PETERS, M. (1997): "A Competitive Distribution of Auctions," *Review of Economic Studies*, 64, 97–124.
- 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.
- RUBINSTEIN, A., AND A. WOLINSKY (1990): "Decentralized Trading, Strategic Behaviour and the Walrasian Outcome," *Review of Economic Studies*, 57(1), 63–78.
57(189), 63–79.

- 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 (1989a): "The Rate of Convergence to Efficiency in the Buyer's bid Double Auction as the Market Becomes Large," *Review of Economic Studies*, 56, 477–498.
- SATTERTHWAITE, M. A., AND S. WILLIAMS (1989b): "Bilateral Trade with the Sealed Bid k-Double Auction: Existence and Efficiency," *Journal of Economic Theory*, 48, 107–133.
- SMITH, V. (1962): "An Experimental Study of Competitive Market Behavior," *Journal of Political Economy*.
- 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>.