



Internet auctions with many traders

Michael Peters^a, Sergei Severinov^{b,*,1}

^a*Department of Economics, University of British Columbia, Canada*

^b*Fuqua School of Business, Duke University, USA*

Received 21 January 2003; final version received 14 April 2005

Available online 24 June 2005

Abstract

We study a multi-unit auction environment similar to eBay. Sellers, each with a single unit of a homogeneous good, set reserve prices at their own second-price auctions. Each buyer has private value for the good and wishes to acquire a single unit. Buyers can bid as often as they like and move between auctions. We characterize a perfect Bayesian equilibrium for this decentralized dynamic mechanism in which, conditional on reserve prices, an efficient set of trades occurs at a uniform price. In a large but finite market, the sellers set reserve prices equal to their true costs under a very mild distributional assumption, so ex post efficiency is achieved. Buyers' strategies in this equilibrium are simple and do not depend on their beliefs about other buyers' valuations, or the number of buyers and sellers.

© 2005 Elsevier Inc. All rights reserved.

JEL classification: C73; C78; D44

Keywords: Decentralization; Trading mechanisms; Auctions; Internet

1. Introduction

An important class of allocation problems involves multilateral exchanges where many sellers and buyers meet to trade. Mechanism design offers a good deal of advice how to organize trade in such environments. When efficiency is the primary concern, some modification of the Vickrey–Clarke–Groves (VCG) mechanism is a natural candidate. The VCG mechanism has desirable incentive properties and allows to implement an efficient allo-

* Corresponding author.

E-mail addresses: peters@econ.ubc.ca (M. Peters), sseverin@duke.edu (S. Severinov).

¹ This paper was completed when this author was visiting the Hoover Institution, Stanford University.

cation in many environments, including the heterogeneous goods case. In an environment where a single seller allocates the goods to privately informed buyers, the VCG mechanism maximizes the seller's revenue within the class of ex post efficient mechanisms [23,10]. The VCG mechanism runs an expected deficit when all traders have private information. For this reason, it might not be feasible.

When budget balance and individual rationality constraints must be satisfied, double auctions are a natural trading mechanism. In the homogeneous goods case, double auctions are *interim incentive efficient* when there are sufficiently many buyers and sellers [21], and become *ex post efficient* quickly as the number of buyers and sellers increases [19]. In a *seller's offer* double auction,² the buyers' (but not generally sellers') payoffs are the same as in Vickrey auction.

A natural generalization of the seller's offer double auction in the environments where goods are not identical³ is a mechanism where buyers describe their willingness to pay for every good, while sellers announce their ask prices. An auctioneer then computes an allocation that is efficient given the announced preferences, and sets the price in every trade equal to the *buyer's* Vickrey price. It has not yet been shown whether the desirable efficiency properties of double auctions generalize to this mechanism, but it seems likely that they would.

Finally, there are fairly simple algorithmic procedures that can be used to determine the Vickrey prices and allocations. One such procedure for the case where each buyer wants only a single good is described in [18]. Ausubel and Milgrom [2] study a dynamic proxy bidding auction that leads to a Vickrey allocation in a Nash equilibrium within a setting where buyers have demands for bundles of goods.

An important feature of all these mechanisms is their reliance on the centralized processing of demand and supply information. Specifically, both in double auctions and in Vickrey mechanisms buyers and sellers send messages to a center which uses this information to compute an array of trades and prices, and then sends back the appropriate instructions to buyers and sellers.

For reasons that are not yet completely clear, practise and theory usually diverge at this point—especially in the case of popular internet auction markets.⁴ In eBay, Ubid, Amazon, or even Sotheby's auction mechanisms, market demand and supply information is not collected or aggregated at all. The auction house sets the overall trading rules, but it neither calculates market clearing prices, nor sends instructions to buyers and sellers regarding the trades which they have to execute. Instead, trading is organized in a way that

² If there are m buyers and n sellers, then in a "seller's offer" double auction the trading price is set equal to the m th lowest value (from the bottom) among buyers' bids and sellers' asks.

³ Bajari and Hortacsu [3] describe an internet market for coins in which the objects being traded are obviously not identical, but where the market otherwise looks very competitive. A similar example of an internet market for antiques is described in [17]. Lucking-Reiley [11] observes that on eBay and other sites many sellers of the same good often run their auctions simultaneously.

⁴ Internet exchanges continue to grow rapidly despite the end of the dot-com boom. The volume of trade at business-to-business internet marketplaces has grown from \$5 bln. in 1999 to more than \$100 bln. in 2001 (see [22,13]), while in consumer auction markets the volume of exchange was projected to grow from \$ 12 bln. in 2002 to \$54 bln. in 2007 (see [8]). The largest consumer-oriented site—eBay—had reported 56.1 million active users in December 2004. eBay's revenue in 2004 reached \$3.27 bln.

looks more like a competing auction market [16,12,15] than a double auction or Vickrey mechanism. This is so despite the fact that some segments in these markets appear to be ideal candidates for a simple double auction.⁵ According to the rules of these mechanisms, a seller is allowed to communicate with some set of buyers before choosing a partner and setting her price.⁶ An important principle appears to be that the price charged by a seller and the identity of the seller's trading partner can depend only on the messages sent by interested buyers, but cannot depend on information about other sellers (such as their reserve prices), or on messages sent by buyers to other sellers—properties that are fundamentally at odds with the Vickrey approach.⁷

Intuitively, we can regard these trading mechanisms as *decentralized*, in contrast to the mechanism with centralized processing of demand and supply information. It is hard to provide an encompassing and general definition of a decentralized mechanism. Below we offer a definition which can be applied in this and similar environments where several sellers offer their goods for sale to a group of buyers.

Any trading mechanism can be formally viewed as a mapping that converts the messages sent by market participants into allocations for the entire market. We will say that a mechanism is *decentralized* if each buyer's message is a collection of separate messages sent to different sellers—one for each seller, and all actions of a particular seller and her final allocation are independent of messages that the buyers send to other sellers.⁸ In essence, our definition simply says that the grand mechanism should be decomposable into a collection of mechanisms run by individual sellers.

In this paper we do not try to explain what factors make decentralized mechanisms attractive to buyers and sellers. Instead, we simply take for granted the existence of such decentralized trading institutions, and focus on the incentive and coordination issues that this institutional design generates.

Specifically, consider the following two aspects. First, when sellers run separate auctions, buyers can independently communicate with many different sellers and thus manipulate trading outcomes at a number of them. One implication of this is that the incentive properties of any seller's mechanism will typically depend on the mechanisms offered by the other sellers. Second, an efficient outcome cannot be attained in a decentralized mechanism unless the exchange of messages between buyers and sellers was truly dynamic. At the same time, a dynamic procedure gives the traders an opportunity for manipulation at each stage.

The goal of this paper is to demonstrate that despite potential obstacles, there is a decentralized trading mechanism for the multilateral exchange which possesses a *perfect* Bayesian equilibrium supporting ex post efficient trade at Vickrey prices. In our mechanism sellers

⁵ Atsushi Kajii has suggested that video game software, which does not depreciate or vary in quality, could easily be sold in a double auction. Instead, a number of manufacturers offer the games for sale on eBay with each unit being sold in a separate auction.

⁶ The institutions themselves clearly constrain the set of messages that buyers and sellers are allowed to exchange. On eBay, a buyer can only submit bids or accept prices, but cannot negotiate with the seller.

⁷ The price that a buyer pays in a Vickrey mechanism is chosen so that (s)he obtains exactly the marginal surplus that (s)he generates. This marginal surplus depends on the costs of all sellers and on the valuations of all buyers.

⁸ Again referring to eBay, a buyer's message is a bid or a set of bids at every seller. A particular seller's price and the identity of the buyer who she trades with can only depend on the bids that the buyers submit to her.

independently run ascending second-price auctions where they are free to set reserve prices, and buyers bid in multiple rounds. If their bids are not successful, the buyers can adjust them and move between sellers costlessly. The bidding procedure adapts the process analyzed in the literature on competing auctions (see, for example, [16]) by making it dynamic. This mechanism is similar (but different in some important ways) to the one used by eBay.

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. 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 slowly as possible (as long as the bid is below his valuation). The only two pieces of information that a bidder needs are the standing bid, i.e. the current price in the auction, and whether the standing bid has changed since the last change of the winning bidder. In most internet auctions, the standing bid and the identity of the winning bidder are typically published at all times. So to follow our rule, a bidder only has to monitor the changes in the announced data. She does not need to directly observe whether and when other bidders submit their bids.

Interestingly, following this rule constitutes a perfect Bayesian equilibrium in the bidding process independently of buyers' beliefs about other buyers' valuations, and even the number of other buyers. The outcome of this equilibrium is efficient provided that sellers set their reserve prices equal to their true costs. So we proceed to show that it is an equilibrium for all sellers to set reserve prices equal to their true costs when the number of traders participating in the market is sufficiently large (but still finite). The remarkable part of our results is that the outcome of the bidding process is efficient and sequentially rational (i.e. optimal at every information set given the traders' beliefs and their strategies), yet *looks* very much like a simple algorithmic price adjustment procedure.

Our results are related to the work on the price adjustment procedures described, for example, by Roth and Sotomayor [18], and on proxy bidding auctions by Ausubel and Milgrom [2]. There are at least two aspects that distinguish our approach from their work. First, we insist that buyers follow sequentially rational strategies at each information set in the bidding game, so the final outcome of our bidding procedure constitutes a perfect Bayesian equilibrium.⁹ Second, many different sellers participate in our mechanism and we analyze their strategic incentives in it, while Roth and Sotomayor [18] and Ausubel and Milgrom [2] do not deal with the issue of the seller's incentives at all.

The results in the second part of the paper regarding the sellers' incentive to set their reserve prices at or above their costs, are closely related to those in the literature on large double auctions, e.g. [20,19], and more recently [6]. We show that our decentralized mechanism is outcome equivalent to the seller's offer double auction. In particular, the sellers who set their reserve prices in independent auctions in our mechanism are exactly in the same strategic situation that they face in this double auction. Satterthwaite and Williams [20] and Rustichini et al. [19] demonstrate that sellers will set reserve prices near their true costs

⁹ Ausubel and Milgrom [2] consider an environment with heterogeneous goods and more general buyer demands. In their auction, buyers submit bids through a proxy bidder in a manner that causes prices to rise in much the same way that they do here. The proxy bidding strategies constitute a Nash equilibrium, yet there is no presumption that the bids submitted by the proxy bidder along the price adjustment path are sequentially rational for the real bidder.

when there are many traders in a double auction, and the costs and valuations are distributed independently. They do not prove existence of equilibria in double auctions. This problem is resolved in [9] who establish the existence of an equilibrium with trade in private value double auctions without requiring independence. Cripps and Swinkels [6] show that in all equilibria where trade occurs with a positive probability, prices must be close to competitive prices when there are enough traders—again without requiring independence.

We obtain somewhat stronger results by considering an environment where sellers and buyers bid on a finite grid. We are able to show the existence of a fully ex post efficient equilibrium when the number of buyers and sellers is large but finite. This result holds under a very mild distributional assumption that each trader's valuation/cost takes every value on the grid with a positive probability.¹⁰

To summarize, the overall contribution of the paper is two-fold. First, we construct a decentralized fully dynamic bidding mechanism and demonstrate the existence of an ex post efficient perfect Bayesian equilibrium in it, without assuming that valuations are independently distributed. Second, we strengthen the available convergence results for large auctions by demonstrating that the sellers will post prices equal to their costs under a very mild distributional assumption that a trader's valuation takes each possible value with a positive probability.

2. The model

There are n sellers and m buyers trading in a market. 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 $\mathcal{D} \equiv \{\underline{p}, \underline{p} + d, \underline{p} + 2d, \dots, \bar{p}\}$ that has a step size $d > 0$. A buyer with valuation b who wins a single auction at a price p gets surplus $b - p$. A buyer who wins more than one auction gets no additional utility from the additional units of output (so his payoff will decrease because he has to pay for the additional units). A seller with cost c who sells at price p gets surplus $p - c$.

Let $F^{m,n}(\cdot)$ denote the probability distribution from which the array of buyers' valuations and the sellers' costs are drawn in the market with m sellers and n buyers. Our results on equilibrium bidding behavior by buyers are independent of $F^{m,n}(\cdot)$ and the buyers' beliefs about it. In the analysis of the seller's part of the game we assume only that, given any number of buyers and sellers and the array of valuations and costs of other traders, a particular buyer's valuation or a seller's cost take each value on the grid with a probability that is bounded above zero. Thus, our results apply in both correlated and independent private value environments.

Trade is organized in the following way. At first, sellers simultaneously announce reserve prices in their auctions. Thereafter buyers arrive sequentially. When a new buyer arrives, he is given an opportunity to submit one or more bids at whichever of the sellers' auctions

¹⁰ A similar distributional assumption is used by Dekel and Wolinsky [7] to establish the convergence result in a *single-object* first-price auction with a discrete grid of bids and values. They show that each buyer will choose the bid closest from below to her true value when the number of buyers is large. The main difference of our case is that we allow both the number of buyers and the number of objects (sellers) to increase at the same rate.

he likes best. Buyers are required to submit bids in the grid \mathcal{D} .¹¹ 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 assume that, as in most on-line auctions, the identity of the high (winning) bidder is made public by each seller,¹² though the high bid itself is not revealed (although it can be inferred if the standing bid is raised later). 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 winning bidder. The standing bid in this case is equal to the high bid.

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

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 buyers' bidding game 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 sequentially rational strategy in a perfect Bayesian equilibrium for buyers to bid their true valuations when they start bidding. We illustrate this point in some detail with an example, since it provides motivation for the bidding rule that we study below.

To see the argument, suppose that there are two buyers 1 and 4 with true valuations b_1 and b_4 and two sellers who announce reserve prices s_1 and s_2 . Let $b_1 > s_2 > s_1$. Visually, consider the data in Fig. 1, but ignore points b_2 and b_3 and allow b_4 , which is not shown in the figure, to take all possible values. Suppose that buyer 1 with valuation b_1 enters first and expects buyer 4 to bid his true valuation b_4 at one of the sellers.

¹¹ This assumption is natural in view of our interpretation that the grid on which the traders' valuations are distributed is determined by the minimal monetary unit.

¹² As 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.

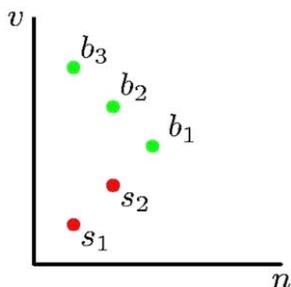


Fig. 1.

We will show that buyer 1 can earn a higher expected payoff if, instead of bidding 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 with seller 2, then it is optimal for buyer 4 to bid with seller 1 who has a lower reserve price. No matter what bid buyer 4 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 1, because there is a positive probability that buyer 4's valuation is below s_2 . Even if buyer 4's valuation is above s_2 and she bids it at seller 1, the worst that can happen to buyer 1 is that she bids b_1 at seller 2 and gets the good at s_2 .

Next, suppose that buyer 1 bids his valuation with seller 1. When buyer 4 enters, the standing bid with seller 1 is still equal to s_1 , because seller 1 has yet to receive a second bid. If buyer 4 bids at seller 2, then the bidding ends and buyer 1 trades with seller 1 at price s_1 . The same outcome would occur if buyer 1 bid s_2 , and not her valuation, initially. If buyer 4 bids her valuation at seller 1, then three cases are possible.

Case 1: $b_4 \leq s_2$. Then buyer 1 will trade at price equal to $\max\{s_1, b_4\}$. The same outcome will occur if initially buyer 1 bids s_2 , instead of b_1 , at seller 1.

Case 2: $b_4 > b_1$. 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.

Case 3: $b_4 \in (s_2, b_1]$ (which occurs with a strictly positive probability). If buyer 1 bids b_1 at the start, then she will trade with seller 1 at price equal to b_4 . However, if buyer 1 starts by bidding s_2 , she will be displaced as high bidder at seller 1. She can then bid at seller 2 and trade with her at price s_2 . Thus, bidding s_2 ensures that buyer 1 never has to pay a price above s_2 and improves her expected payoff.

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 valuation, have an option to bid again elsewhere. Unfortunately, this option is not sufficient to guarantee that, conditional on the sellers' reserve prices, the efficient trades are carried out. We will demonstrate below that the dynamic bidding game has multiple equilibria some of which are inefficient. 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 $\mathbf{c} = \{c_1, \dots, c_n\}$ ($\mathbf{s} = \{s_1, \dots, s_n\}$) be the vector of sellers' costs (reserve prices). Let $v_m(\mathbf{b}, \mathbf{s})$ be the m th lowest (from the bottom) valuation in the vector (\mathbf{b}, \mathbf{s}) . When the argument of this function is clear from the context, we will simply use the notation v_m .

If the sellers' reserve prices are equal to their true costs, then an 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 efficiency implies that m traders who end up *without* the good have the lowest valuations or costs. In competitive terms, the price v_m is the *lowest* price that clears the market given the demand and supply schedules generated by the buyers' valuations and the sellers' reserve prices. It is also easy to see that v_m is a 'Vickrey' price ensuring that every buyer who trades gets an amount of surplus equal to the difference between the maximal *apparent* gains to trade between all buyers and sellers and the maximal apparent gains to trade when buyer i is left out of the allocation, under the given array of the buyers' valuations and the sellers reserve prices.

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

Definition 1. The symmetric strategy σ^* is defined as follows. When 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 who has this lowest standing bid. The bid should be equal to the lowest value on the grid that exceeds this lowest standing bid;
- (c) otherwise, if more than one seller has the lowest standing bid, the buyer should submit the same bid as in (b) but choose among the sellers with the lowest standing bid according to the following rule: select with equal probability each of the sellers where the standing bid has changed since the last change of the winning bidder and sellers who have not yet received a bid. If there are no sellers in either of these two categories, a bidder should randomly choose between all sellers whose standing bid is the lowest.

According to strategy σ^* , a buyer should always bid at one of the sellers with the lowest standing bid. Then active buyers will 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). Active buyers will then continue to bid 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) which tells the buyer how to choose among the sellers with the lowest standing bids. In essence, it allows a buyer to identify which of these sellers have lower winning bids (recall that winning bids are unobservable). Bidding only at such sellers is optimal for the buyer, because then with some probability the buyer will trade at a lower price. At the same time, such bidding ensures that standing bids rise more slowly and eliminates the possibility of price dispersion.

To understand why rule (c) allows to identify the sellers with the lowest winning bids when all buyers use strategy σ^* , note that a new winning bid is always strictly above the standing bid (by one grid point) and displaces the previous winning bidder without changing the standing bid. In contrast, an unsuccessful bid causes the standing bid to rise without displacing the high bidder. Thus, when two standing bids are the same, the one where the standing bid has changed after a change in the identity of the winning bidder will have a lower winning bid, and this is where the buyer should bid according to rule (c).

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 Fig. 1, where the valuations and reserve prices of three buyers and two sellers are shown. This example is designed just to illustrate the path of the bidding game when the buyers use strategies σ^* , 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_1, b_2, b_3 . Buyers enter in order of their indexes, so buyer 1 with valuation b_1 enters first, and buyer 3 enters last. According to strategy σ^* , when buyer 1 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 successful, but it will cause the standing bid with seller 1 to rise to s_2 . Unsuccessful bids always change a seller's standing bid, since buyers must submit bids above the current standing bid. As buyer 1 was the first to bid s_2 , he remains a high bidder and passes by (a), so buyer 2 has a chance to submit a new bid. Now both sellers have the same standing bid. The standing bid with seller 1 has changed after the winning bidder has changed, and seller 2 has not yet received a bid. The lowest value that exceeds the standing bid s_2 is b_1 , so by (c) buyer 2 will bid b_1 with equal probability at seller 1 (case A) or seller 2 (case B). In either case, buyer 2 will become a high bidder but his bid will not affect either of the standing bids. Note that a winning bid never changes the seller's standing bid.

Consider case A first. Since buyer 2 displaces buyer 1, buyer 1 immediately submits a bid equal to b_1 with seller 2 by rule (c), since seller 2 has yet to receive a bid. This bid is successful, but does not affect seller 2's standing bid which remain equal to s_2 . When buyer 3 enters, he finds that at either seller the standing bid has not changed since the last change of the high bidder, so he will bid b_1 at one seller, and then at the other in random order. Neither bid by buyer 3 will be successful, but the standing bids at both sellers will rise to b_1 . Then buyer 3 will bid b_2 choosing randomly between the two sellers (since at each seller the standing bid has risen after the high bidder has changed). If buyer 3 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 high bidder there. Bidding will then stop, and buyers 2 and 3 will trade at price b_1 . If buyer 3 bids b_2 first with seller 2, his bid will be successful and all bidding will stop.

Now consider case B. After buyer 2 becomes a high bidder with seller 2, buyer 3 enters and submits bid b_1 at seller 1 displacing bidder 1 as a high bidder without raising the standing bid above s_2 . Buyer 1 will then submit bid b_1 at sellers 1 and 2 in random order, which will raise standing bids at both sellers to b_1 , but bidder 1 will not become a high bidder. The bidding will then stop and buyers 2 and 3 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, efficient trades occur and both sellers trade at a uniform price equal to buyer 1's valuation. The randomness on the path generated by σ^* makes possible different pairwise matching combinations between buyers and sellers. 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^* make it an equilibrium choice, as the following theorem demonstrates:

Theorem 1. *It is a perfect Bayesian equilibrium of the bidding game for all buyers to use the strategy σ^* . For each array of valuations and reserve prices (\mathbf{b}, \mathbf{s}) , each buyer whose valuation is above v_m trades with some seller whose reserve price is no larger than v_m . Each seller whose reserve price is below v_m trades for sure with some buyer whose valuation is at least v_m . All trades occur at price v_m .*

Proof. See the appendix. \square

The central implication of Theorem 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 lowest valuations and reserve prices. Thus, apparent gains from trade are maximized.

The strategy σ^* contains the same set of rules on and off the equilibrium path. But although the equilibrium path is relatively simple, the proof that playing σ^* constitutes a perfect Bayesian equilibrium is sufficiently complex. Perfect Bayesian equilibrium requires σ^* to be sequentially rational, i.e. for each bidder σ^* must be a best reply to the other bidders' play of σ^* at every information set, even the ones which are 'far away' from the equilibrium path and can be reached only after multiple deviations from σ^* by several players. Showing this turns out to be a challenging task due to the complexity of the dynamic bidding game involving multiple auctions and the buyers' switching between them.

Our proof works as follows. We fix an arbitrary information set in the bidding game and consider the continuation game starting from it. First, we characterize the outcome of this continuation game when all players follow σ^* . This part of the proof also provides a characterization of the candidate equilibrium outcome obtained when all players always follow σ^* from the start of the game. Then we demonstrate that an arbitrarily chosen bidder i cannot improve her payoff in this continuation game by a unilateral deviation from σ^* . Specifically, we show the following. If buyer i submits a bid which is higher than what is prescribed by σ^* and/or submits multiple high bids whereas all other buyers follow σ^* , then the trading price would be at least as high and the expected number of units that i would purchase would be at least as large as when she does not deviate from σ^* in the continuation game. Since buyer i purchases the desired number of units at the trading price that prevails when all buyers follow σ^* in the continuation, such deviation cannot be profitable.¹³ At the same time, a buyer who at any stage 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.

Note that neither the description of strategy σ^* , nor the proof that σ^* is a best response depend in any way on the distribution of valuations, or the number of buyers and sellers in

¹³ By the one-deviation property, it is sufficient to consider buyer i 's strategies that involve a deviation from σ^* at a single information set in the continuation game.

an auction. Our proof does impose some restrictions on buyers' beliefs in off-equilibrium continuation games. First, when buyers observe data that are inconsistent with all buyers' using the strategy σ^* , they believe that the underlying deviation is 'minimal' in an appropriate sense. In particular, a bid submitted by a deviator is believed to be as close as possible to the bid prescribed by σ^* provided that it can still produce the observed data.

Second, all buyers believe that the valuation of the high bidder in any auction is at least as high as the standing bid in that auction. This imposes a restriction in the case where the buyers are certain about each other's valuation at the beginning of the bidding process or the distributions of the valuations do not have full support. Consider the former case, for example. It is possible that some buyer might deviate and become a high bidder with a seller whose standing bid is above that buyer's valuation. We require that buyers revise their initial beliefs in this case, and believe that the deviating buyer's valuation is at least as high as her bid. However, this belief rule, unlike the first one, is not necessary for our equilibrium and any alternative one would do as well, while the first rule does not involve beliefs about valuations. Therefore, we can say that σ^* constitutes an 'ex post' equilibrium, i.e. it remains an equilibrium even if buyers know each other's valuations. We regard this as an important property confirming the robustness of our equilibrium.

Another important property of the equilibrium where buyers use σ^* is the absence of price dispersion in the final outcome: all trades are executed at the same price. As explained above, rule (c) in σ^* plays a crucial role in this.

The efficiency in the bidding game can be attributed to the rules of our mechanism which, in our interpretation, are designed by a large institution (internet exchange) in the context of which individual sellers offer their goods. It is important to understand to what degree our efficiency results are robust to changes in these rules. Addressing this issue in full generality is outside the scope of this paper. Still, it is easy to see that the non-observability of the high bids combined with our restriction on beliefs off the equilibrium path—we insist on the minimal revision of equilibrium beliefs following a detected deviation—ensure that preemptive bidding does not occur in our mechanism.

On the other hand, if the high bids were observable, then the game would have an equilibrium in which the first buyer in order of entry submits the highest possible bid \bar{p} at the seller with the lowest reserve price, the second buyer in order of entry submits the same bid \bar{p} at the seller with the second lowest reserve price, and so on. No other buyer would bid at a seller who has received the bid \bar{p} . Thus, buyers who are among the first to enter would be 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.

The bidding problem that we consider is similar to the matching environment discussed in [18]. Their environment differs from ours in that buyers may have different valuations for the goods offered by different sellers, while in our case the goods are homogeneous, and we fully analyze the traders incentives. Consider any stage of our bidding process and suppose that all buyers follow strategy σ^* . If the number of buyers willing to pay strictly more than the lowest standing bid is greater than the number of sellers with this standing bid, then the set of goods offered at this lowest standing bid constitutes what Roth and Sotomayor call an *overdemanded set*. The subsequent bidding will raise each of the standing bids of

sellers in that set by one grid point. In this sense our bidding rule implements the algorithm discussed by Roth and Sotomayor. The content of Theorem 1 is to show that there is a perfect Bayesian equilibrium for which the algorithm is not subject to strategic manipulation at any time during the course of the procedure. What makes the argument difficult is the fact that we have to deal with situations that could never occur in the bidding process under the simple application of the Roth–Sotomayor algorithm.

For each array (\mathbf{b}, \mathbf{s}) of buyers' valuations and sellers' reserve prices, Theorem 1 uniquely specifies the price at which all trades will be executed. Precisely, the trading price will be equal to $v_m(\mathbf{b}, \mathbf{s})$ which is either the highest valuation among buyers who fail to trade, or the highest reserve price among sellers who trade, depending on the actual array of valuations and reserve prices. Buyers whose valuations strictly exceed v_m and sellers whose reserve prices are strictly below v_m will trade for sure. At the same time, the randomization on the equilibrium path (buyers randomize between the sellers among whom they are indifferent) implies that the outcome for traders whose valuations or reserve prices are exactly equal to $v_m(\mathbf{b}, \mathbf{s})$ may be random: they may or may not trade.

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 any array (\mathbf{b}, \mathbf{s}) , 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 and sellers in N_1 will trade for sure, and that buyers in M_1 and sellers from N_3 will not trade.

Corollary 1.¹⁴ *If all buyers use the strategy σ^* in the bidding game, then the number of sellers with reserve prices equal to v_m who trade is between*

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

while the number of buyers with valuations equal to v_m who trade is between:

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

Of course, the efficiency in the buyers' game does not guarantee that the equilibrium outcome of the mechanism will be efficient, since the sellers may set reserve prices that are different from their true costs. Before we turn to the issue of the sellers' incentives in the next section, it is worth pointing out that our bidding game has multiple equilibria. The following example illustrates another plausible, yet inefficient equilibrium.

In Fig. 2, given sellers' reserve prices s_1 and s_2 the efficient outcome is for buyer 2 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 enters last, because he

¹⁴ The Proof of the Corollary is available at http://microeconomics.ca/sergei_severinov/proofsinternetauc.pdf.

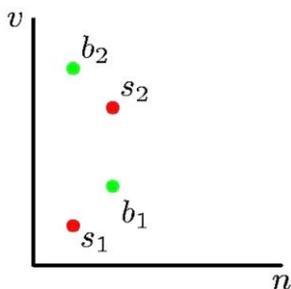


Fig. 2.

believes (correctly) that the buyer who has already entered has bid s_2 at seller 1. Therefore, there is a perfect Bayesian equilibrium in which both buyers use this strategy. Given the data in Fig. 2, in this equilibrium buyer 1 will trade with seller 1, and buyer 2 with a higher valuation will trade with seller 2. Note that the existence of this inefficient equilibrium is invariant to the order of moves: if we reverse the order of entry and let buyer 2 enter first, inefficiency would obtain when $b_1 > s_2 > b_2 > s_1$.

The example demonstrates that our mechanism does not support efficient trade as a unique outcome. Nevertheless, we believe that it is unrealistic to hope for simple mechanisms that will uniquely support efficient trade. Typically, complex and unrealistic mechanisms need to be constructed to fully implement efficient allocations. We leave this problem for future research. Our goal is to demonstrate that efficiency could be attained with ‘plausible’ looking indirect mechanisms in strategies robust to perturbations in the traders’ beliefs about each other.

4. Optimal reserve prices in a large market

The results of the previous section indicate that the buyers’ equilibrium strategies guarantee that an efficient set of trades occurs when sellers set reserve prices equal to their true valuations. Yet, the outcome will not necessarily be efficient if sellers set reserve prices that are different from their true costs. So we turn to an examination of sellers’ behavior in this section.

It may not be optimal for a seller to set a reserve price equal to her true valuation, because her reserve price may in some cases affect the trading price. For example, consider the situation depicted in Fig. 1 without buyer 1 (the one with the lowest valuation) and let buyer 2’s valuation b_2 still be undetermined. From this 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. In fact, there is a close link between our decentralized mechanism and the “seller’s offer” double auction where the trading price is set equal to the m th lowest value (from the bottom) among

buyers' bids and sellers' asks. Satterthwaite and Williams [20] show that a buyer in this double auction will optimally bid his true valuations. So, the trading price in the "seller's offer double auction" is also equal to v_m . Since in our decentralized market and in the "seller's offer double auction" all sellers set their reserve prices simultaneously, the payoff that a seller gets after posting a particular reserve price is the same in both mechanisms for every array of buyers' valuations and other sellers' reserve prices. So, from the seller's point of view they are strategically equivalent. Consequently, the sets of equilibrium outcomes and associated prices must also be the same in the two mechanisms (provided that the buyers follow the strategy σ^* in our mechanism).

Our next result establishes that in our framework with a finite grid of valuations, an ex post efficient equilibrium in which sellers post reserve prices equal to their true costs exists when the number of traders is sufficiently large, but finite. This result holds under a very mild distributional assumption that each trader's valuation/cost takes every value in the grid \mathcal{D} with a positive probability which is bounded above zero. Thus, we extend the results of Rustichini et al. [19] who show that the sellers' optimal ask prices and buyers' optimal bids converge to their true costs and valuations in a double auction when the costs and valuations are independently and continuously distributed over an interval.¹⁵

Slightly modifying the notation used in the previous section, let \mathbf{b}^m (\mathbf{b}_{-i}^m) denote an array of the realized valuations of all buyers (all buyers other than i) in an m -buyer market. Similarly, let \mathbf{c}^n (\mathbf{c}_{-j}^n) denote an array of the realized costs of all sellers (all sellers other than j) in an n -seller market. Further, let $f^i(p|\mathbf{b}_{-i}^m, \mathbf{c}^n)$ ($g^j(p|\mathbf{b}^m, \mathbf{c}_{-j}^n)$) denote the conditional probability that buyer i 's (seller j 's) valuation (cost) is equal to p given the profile $\mathbf{b}_{-i}^m, \mathbf{c}^n$ of other buyers' and sellers' valuations and costs. Finally, let $f(p)^{m,n} = \min_{1 \leq i \leq m, \mathbf{b}_{-i}^m, \mathbf{c}^n} f^i(p|\mathbf{b}_{-i}^m, \mathbf{c}^n)$, and $g(p)^{m,n} = \min_{1 \leq j \leq n, \mathbf{b}^m, \mathbf{c}_{-j}^n} g^j(p|\mathbf{b}^m, \mathbf{c}_{-j}^n)$. Note that the minimum is taken across all buyers (sellers) and across all possible profiles of other buyers' valuations and sellers' costs. These minima exist because the number of buyers and sellers and the number of feasible valuation/cost profiles are finite. Our main assumption is as follows:

Assumption 1. $\forall p \in \mathcal{D} \exists f(p) > 0, g(p) > 0$ s.t. $f^{m,n}(p) \geq f(p)$ and $g^{m,n}(p) \geq g(p)$ $\forall m, n$.

Assumption 1 is satisfied in many well-known environments. For example, it holds when a buyer's valuation and a seller's cost are conditionally independent given the realization of some random variable, and all conditional distributions have a full support.

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., $m = kn$. 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 the sellers when the number of traders in the market is sufficiently large.

¹⁵ Rustichini et al. [19] have to assume the existence of an equilibrium where trade occurs with a positive probability, while we directly establish the existence of an equilibrium in our framework.

Theorem 2. *Suppose that Assumption 1 holds, and that $f(p)$, $g(p)$, ($p \in \mathcal{D}$), and the number of buyers m and sellers n in the market are common knowledge. If the buyers follow the strategy σ^* , and if m and n are sufficiently large, then it is a perfect Bayesian equilibrium for each seller to set her reserve price equal to her true cost.*

Theorems 1 and 2 together imply that an ex post efficient outcome is attained in our decentralized mechanism when the number of traders is sufficiently large but finite. For this result to hold, the buyers need not have any information about other traders, while the sellers only need to know the number of other traders and the (positive) lower bound on the probability that a buyer's valuation and the seller's cost takes a particular value.

Thus, this paper contributes to the literature that searches for mechanisms that are 'detail-free' i.e. mechanisms whose design does not depend on the distributions of the traders' valuations and the common knowledge of these distributions.

The proof of Theorem 2 demonstrates that seller j with cost $c < p$ obtains a higher expected payoff by setting reserve price equal to $p - d$ rather than p when the number of traders is sufficiently large, and all other sellers set their reserve prices equal to their true costs. To understand this result, consider the tradeoff that a seller faces when she decides whether to set a higher reserve price $p > c$ or cut it to $p - d$.

The cost to a seller of cutting the reserve price is that she may actually trade at price $p - d$, while she would have traded at price p had she set the reserve price at that higher level. In other words, the cost is incurred if a seller trades irrespectively of which of the two reserve prices she posts, but her decision affects the price in the market. The seller gains from cutting her reserve price if this allows her to trade (either at price $p - d$ or p), while she does not trade after posting p .

It is easy to see that the gain outweighs the cost of cutting the reserve price, if the probability that a seller trades at p when she posts price $p - d$ is greater than the probability that she trades at p when she posts price p . Of course, the probability that trade occurs at any particular price goes to zero as the number of traders increases. But when we compare these two probabilities, we can condition on the event that price p is *pivotal*, i.e. the profile of all buyers' valuations and $n - 1$ sellers' costs is such that if the remaining seller j posts reserve price p , then the trading price would be equal to p . If price p is pivotal, then seller j posting reserve price $p - d$ always trades, but the trading price could be either p or $p - d$.

So, we can restate the desired claim as follows: Conditional on the event that price p is pivotal, the probability that a seller's decision to post price p or $p - d$ does not affect the trading price (i.e. the price remains at p) is greater than the probability that seller posting price p trades. The proof establishes that this claim is true when the number of traders in the market is sufficiently large.

To understand the intuition behind this result, suppose that all sellers other than j set their reserve prices equal to their true costs. Let p be pivotal, i.e. equal to the $(m - 1)$ -th lowest value in the array consisting of the valuations of all buyers and the costs of sellers other than j . Given that the number of both buyers and sellers is large, the expected number of sellers with cost p and buyers with valuations p is also large.

The last observation has two implications. First, by Corollary 1, if the seller sets a reserve price equal to the trading price, she will be competing with all other sellers who set this price. Because of this competition a seller posting pivotal price p will often fail to trade.

Second, conditional on the event that p is the $(m - 1)$ -th lowest values in the vector of valuations of all buyers and costs of sellers other than j , the probability that p is also the m -th lowest value in this vector is large. But if both $(m - 1)$ -th and m -th values in this vector are equal to p , then j 's decision to set reserve price at p or $p - d$ does not affect the trading price: it will be equal to p either way. So, conditional on p being pivotal, the probability that seller j cannot affect the trading price by lowering his reserve price from p to $p - d$ is large. In fact, we can show that this probability goes to 1 as the number of traders increases. These observations explain why the desired claim is true.

5. Conclusions

Several remarks are in order about the results of this paper. First, the equilibrium in buyers' bidding game that we describe is not unique. Examples in the paper illustrate that alternative equilibria exist and do not generally guarantee efficient allocations. This is not surprising, since an equilibrium in a decentralized market involves coordination of matching decisions of many buyers and sellers. Coordination problems almost always have multiple equilibria and having to choose among them seems inevitable.¹⁶

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, i.e. robust to perturbations in the traders' beliefs about each other. It also implements an efficient allocation. These properties make it reasonable to believe that this equilibrium is focal, and that traders will coordinate on this equilibrium eventually.

To play strategy σ^* , a buyer needs only three pieces of information: the current profiles of high and standing bids and whether a seller's standing bid has changed after the change of the high bidder. Thus, σ^* can be easily implemented via a software robot similar to the one that is used by eBay, making buyer's bidding costs (e.g. costs of time and attention) negligible. Such software would be extremely simple: it needs to keep track of only three pieces of information per seller and would always use the same rule to compute the bid.

It is natural to expect some discrepancy between our results and empirical observations, since our model does not reproduce all the details of the bidding behavior at the eBay, Ubid, Amazon, and 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 at these sites. Roth and Ockenfels [17] suggest network congestion and unexpected demands by the family as reasons why bidders may not be able to revise a bid as intended.

Nonetheless, we believe that our model does provide some insights into the workings of these institutions and some empirically testable implications. Our theory implies that bidding behavior of buyers in an auction on eBay cannot be determined in isolation. The decision where and what to bid is largely affected by other alternatives available at the time.

¹⁶ Second-price auctions, for example, possess asymmetric equilibria in which bidders do not bid their true valuations. There are multiple equilibria in centralized mechanisms, such as double auctions, as well.

Furthermore, internet auctions provide bidders with more information, and make it easier for them to communicate with many sellers and coordinate their behavior. In particular, patient revision of bids generates signals which the buyers need in a decentralized market to coordinate their bidding. All these factors imply that internet auctions are quite different from the one shot auctions that dominate auction theory.

The bidders' behavior in our model is based on one simple principle: a bidder carefully explores trading opportunities with all sellers before raising her bid with a given seller. An extension of this logic suggests an explanation for observed flurry of active bidding close to the end dates in the eBay auctions. 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 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 in an auction that is about to close may also induce a flurry of bidding at other auctions for similar goods.

Anwar et al. [1] have tested some of the implications of our theory of traders' behavior in competing auctions in the context of their study of computer CPU auctions on eBay. They find that a significant proportion of buyers bid across several competing auctions and usually place bids at auctions with the lowest standing bid. Also, more bidders bid across auctions when the difference of ending times between the auctions decreases. Anwar et al. [1] do not specifically study the effect of competing auctions on the price profile, although in their sample prices are typically not uniform which, in our view, could be due to the fact that auctions do not end simultaneously on eBay.

More empirical work remains to be done on this and other issues arising in competing auctions. Although to this date empirical studies of internet auctions has typically focussed on single-auction environments,¹⁷ we hope that this paper will help to stimulate an interest of empirical researchers towards the study of competing auctions. Understanding traders' behavior in such environments appears to be a fairly important task.

Acknowledgments

We thank participants at numerous conferences and seminars for helpful comments and suggestions. We are responsible for all errors.

Appendix

Proof of Theorem 1. 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

¹⁷ For an insightful survey of the relevant literature, see Bajari and Hortacsu [4].

bids together with the identities of the winning bidders, the history of the standing bids and winning bids, and the order in which the buyers move. There is a one-to-one relationship between the nodes in the game and its states. Precisely, the state of the game is a full description of the corresponding node in the game. Define *public state of the game* as the union of all 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 assign indices to buyers based on the order of entry, with buyer 1 arriving first and having the first opportunity to submit new bid(s) after being outbid, and so on.

At each *information set* where a buyer is called to move, this buyer knows the public state of the game, her own high bids and her bidding history. Information sets are partially ordered: one information set precedes the other if the latter can be reached via some sequence of moves from the former. A *path* of the game is a collection of all information sets such that for any pair of information sets in it one can be reached via some sequence of moves from the other. Along any path, one information set precedes the other iff the standing bid at each seller at the former information set is (weakly) lower than at the latter. We will often refer to this property in the proof. Finally, we will say that the bidding game is at the terminal stage if it has reached an information set (or, equivalently, a state) such that no new bids are submitted on the continuation path from it.

We will establish that the bidding game has a perfect Bayesian equilibrium in which all buyers follow the strategy σ^* on and off the equilibrium path, and their beliefs in every information set on and off the equilibrium path are described by the following rules:

Belief Rule 1: Inference about high bids. If the standing bid at a seller has changed after the last change in the identity of the winning bidder, then the high bid at this seller is equal to the standing bid. If the standing bid has not changed since the last change of the winning bidder, then the high bid is one step d above the standing bid.

Belief Rule 2: A buyer's valuation is at least as large as the maximum of her inferred high bids which she holds or has held at some prior information set, with inference done according to Belief Rule 1.

Belief Rule 3: The posterior beliefs about a buyer's valuation held by other buyers are characterized by a posterior distribution obtained from the prior by conditioning on the event that this buyer's valuation is at least as large as the cutoff level inferred according to Belief Rule 2.¹⁸

Clearly, these beliefs are rational on an equilibrium path where all bidders follow σ^* . 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. Our specification of beliefs allows this. Second, if a seller's standing bid changes more than once without a change in the identity of the winning bidder, then the winning bidder must have deviated from σ^* . In this case, according to Rule 1, the other buyers believe that the deviation is the

¹⁸ Since a buyer also observes the decision of another buyer to pass and not to bid at a certain information set, (s)he can also make an inference that the valuation of the buyer who has passed is no higher than the lowest standing bid at that information set. Our results are invariant to this observability assumption and the associated inference, and so we do not include it in the description of the beliefs.

minimal possible one consistent with the path of the game, i.e. the high bid is equal to the standing bid.

The proof that each buyer playing σ^* together with the specified beliefs constitute a perfect Bayesian equilibrium consists of two parts. First, in Lemmas A.1 and A.2 we fix an arbitrary state Γ and characterize the outcome of the continuation game G_Γ starting from an information set corresponding to Γ , when all buyers use the strategy σ^* . Then in Lemmas A.3–A.5 we show that, given the specified beliefs, no buyer can improve her payoff in G_Γ characterized in Lemmas A.1 and A.2 by a unilateral deviation from σ^* . This implies that strategy σ^* is sequentially rational: given the specified beliefs, it is optimal for each buyer to use the strategy σ^* at every information set on and off the equilibrium path provided that all other buyers also use the strategy σ^* .

In fact, our proof relies only on Belief Rule 1 specifying the beliefs about high bids, but not on Belief Rules 2 and 3. Thus, our result is, in fact, stronger. It shows that following σ^* is an ex post equilibrium. That is, it is optimal for the buyers to follow σ^* even if each of them knew the array of other buyers' valuations.

Let us introduce the necessary notation. Consider any point u on the grid, and let $S_\Gamma(u)$ be the set of sellers whose standing bids in state Γ are u or less. Further, let $S_\Gamma^h(u)$ be the subset of $S_\Gamma(u)$ consisting of sellers who in state Γ have high bids strictly greater than u . Similarly, let $S_\Gamma^0(u)$ be the subset of $S_\Gamma(u)$ consisting of sellers who in state Γ do not have high bids strictly greater than u . Obviously, $S_\Gamma^h(u) \cup S_\Gamma^0(u) = S_\Gamma(u)$.

Next, let $D_\Gamma^a(u)$ be the set of all buyers whose valuations are no less than $u + d$ and who in state Γ do not hold any high bids equal to or greater than $u + d$. A buyer in the set $D_\Gamma^a(u)$ is willing to bid in any state of the continuation game where she does not hold a high bid and the lowest standing bid does not exceed u . Set $d_i(\Gamma, u) = 1$ for buyer $i \in D_\Gamma^a(u)$. Also, let $D_\Gamma^p(u)$ be the set of buyers who hold high bids at sellers from $S_\Gamma^h(u)$. For buyer $i \in D_\Gamma^p(u)$, let $d_i(\Gamma, u)$ be the number of such high bids that i holds at Γ . For buyer $i \notin D_\Gamma^a(u) \cup D_\Gamma^p(u)$ set $d_i(\Gamma, u) = 0$. Define $AD_\Gamma(u) = \sum_i d_i(\Gamma, u) \equiv \#D_\Gamma^a(u) + \sum_{i \in D_\Gamma^p(u)} d_i(\Gamma, u)$ (where $\#$ denotes the cardinality of a set).

If all buyers follow σ^* in the continuation game G_Γ and the lowest standing bid at the terminal stage of this game is equal to u , then all high bids held at sellers $S_\Gamma^h(u)$ in state Γ will survive through the terminal stage, and a buyer from $D_\Gamma^a(u)$ will win one unit at a seller from $S_\Gamma^0(u)$. Thus, intuitively, one can view $AD_\Gamma(u)$ as the minimal market demand and $d_i(\Gamma, u)$ as i 's individual demand at price u for units supplied by sellers $S_\Gamma(u)$. Let $g_i(\Gamma, u)$ be the number of high bids held by buyer i in state Γ with sellers whose standing bids are equal to u .

Finally, set $v_\Gamma = \max\{u \mid AD_\Gamma(u - d) > \#S_\Gamma(u - d)\}$ (or, equivalently, $v_\Gamma = \max\{u \mid \#D_\Gamma^a(u - d) > \#S_\Gamma^0(u - d)\}$ since $\sum_{i \in D_\Gamma^p(u)} d_i(\Gamma, u) = \#S_\Gamma^h(u - d)$) if such u exists and let $v_\Gamma = \bar{d}$ if otherwise (recall that \bar{d} is the highest point on the grid). Observe that v_Γ is the minimal market-clearing price at which the demand from the set of active buyers $D_\Gamma^a(v_\Gamma)$ does not exceed the available supply $\#S_\Gamma^0(v_\Gamma)$. Hence, if all buyers follow σ^* in the continuation game, there would not be sufficient demand to drive the lowest standing bid above v_Γ . Indeed, if the lowest standing bid never exceeds v_Γ , then the buyers who in state Γ hold high bids exceeding v_Γ are locked: they would not be active in the continuation game and will end up buying from the sellers where they hold these bids. For the other

buyers strictly wishing to buy at v_Γ , i.e. those in $D_\Gamma^a(u)$, there would be sufficient supply at this price because $\#D_\Gamma^a(u) \leq \#S_\Gamma^0(v_\Gamma)$.

These observations are at the core of the arguments in the following two lemmas which characterize the outcome of the continuation game G_Γ starting from the information set corresponding to an arbitrarily fixed state Γ , when all buyers use the strategy σ^* . To present them, we need one more piece of notation. Say that a high bid in state Γ belongs to set B_Γ^d if and only if it is strictly greater than v_Γ , it is held at some seller j in $S_\Gamma(v_\Gamma)$, and at the information set where this high bid was submitted the standing bid of seller j was strictly below v_Γ . By definition, a submission of a bid from B_Γ^d constituted a deviation from σ^* . As we will show below, this deviation could be relevant on the continuation path of the bidding game.

Lemma A.1. *Consider any state Γ . If all buyers use σ^* in G_Γ , then no trader will trade at a price below v_Γ . All sellers in $S_\Gamma(v_\Gamma - d)$ will trade.*

Proof. The proof is by contradiction. Suppose that the lowest standing bid v_t at the terminal stage \mathcal{T} is strictly less than v_Γ . By definition $\#D_\Gamma^a(v_t) > \#S_\Gamma^0(v_t)$, so there is a buyer $i \in D_\Gamma^a(v_t)$ who is not a winner at any seller from $S_\Gamma^0(v_t)$. Then there are two possible cases. Case (i): buyer i does not trade at all. But i 's valuation v_i is such that $v_i \geq v_t + d$, so, according to σ^* , i would not pass at stage \mathcal{T} . This contradicts the fact that \mathcal{T} is a terminal stage. Case (ii): buyer i trades with a seller $j \notin S_\Gamma^0(v_t)$. But the high bid at seller j in state Γ or her standing bid (if she did not have a high bid) was strictly above v_t . Hence, in G_Γ i must have placed a high bid of at least $v_t + 2d$ at seller j , which contradicts the assumption that all buyers use σ^* in G_Γ .

The lowest standing bid cannot reach v_Γ unless all sellers from $S_\Gamma(v_\Gamma - d)$ receive at least one bid in the continuation game. So, all sellers from $S_\Gamma(v_\Gamma - d)$ will trade. \square

Lemma A.2. *Suppose that all buyers use σ^* in G_Γ . Then the lowest standing bid at the terminal stage will be equal to v_Γ . New bids in G_Γ will be placed only with sellers in $S_\Gamma(v_\Gamma)$, i.e. those whose standing bids in Γ do not exceed v_Γ . Buyer i wins at least $d_i(\Gamma, v_\Gamma)$ units held by these sellers and pays v_Γ for each of them, except for the units of the sellers with whom i holds high bids in B_Γ^d : i will purchase these units either at price v_Γ or $v_\Gamma + d$.*

Proof. First, let us show that the lowest standing bid does not exceed v_Γ at any information set on the continuation path, up to and including the terminal stage. The argument is by contradiction. Suppose that S' is the last information set where the lowest standing bid is equal to v_Γ and the terminal stage is not reached at S' . As σ^* requires that buyers submit bids only at sellers with the lowest standing bid, at information set S' the standing bids of all but one sellers are at least $v_\Gamma + d$ and the standing bid at the remaining seller is equal to v_Γ . Since the next bid at this seller causes her standing bid to rise to $v_\Gamma + d$, the high bid at this seller must be strictly above v_Γ .

Since the lowest standing bid at information set S' is equal to v_Γ , no buyer submits a bid exceeding $v_\Gamma + d$ at any information set preceding S' in G_Γ . Therefore, any seller outside the set $S_\Gamma^0(v_\Gamma)$ has the same high bidder as in Γ . Consequently, only buyers in $D_\Gamma^a(v_\Gamma)$ can be high bidders at sellers from $S_\Gamma^0(v_\Gamma)$. But by definition $\#S_\Gamma^0(v_\Gamma) \geq \#D_\Gamma^a(v_\Gamma)$. Therefore,

there is no buyer at information set \mathcal{S}' whose valuation is at least $v_\Gamma + d$ and who does not hold a high bid. So, no buyer could submit a bid $v_\Gamma + d$ at \mathcal{S}' without deviating from σ^* . In combination with Lemma A.1 this implies that the lowest standing bid at the terminal stage is v_Γ .

Given that all buyers use the strategy σ^* , this result immediately implies the following: (i) a buyer with valuation strictly exceeding v_Γ will trade; (ii) any high bid exceeding v_Γ in state Γ remains a high bid at the terminal stage. Therefore, buyer i will win at least $d_i(\Gamma, v_\Gamma)$ units from sellers in $S_\Gamma(v_\Gamma)$. Additionally, $\#S_\Gamma^0(v_\Gamma) - \#D_\Gamma^a(v_\Gamma)$ units held by sellers from $S_\Gamma^0(v_\Gamma)$ will either be sold to buyers who have valuations equal to v_Γ or who in state Γ hold high bids equal to v_Γ at these sellers, or may remain unsold if their reserve prices are equal to v_Γ ; (iii) on the continuation path from Γ , new bids will be placed only at sellers from $S_\Gamma(v_\Gamma)$, because the standing bids of all other sellers exceed v_Γ in state Γ .

Finally, let us show that the standing bid of a seller in $S_\Gamma(v_\Gamma)$ may reach $v_\Gamma + d$ only if in state Γ this seller has a high bid which belongs to B_Γ^d . Suppose that buyer z 's holds a high bid with seller s_z belongs to B_Γ^d and in information set \mathcal{S}_1 the lowest standing bid and s_z 's standing bid are equal to v_Γ . Then by definition of B_Γ^d , s_z 's standing bid must have changed after z has become a high bidder. So, if a buyer $i \in D_\Gamma^a(v_\Gamma)$ is not a high bidder and moves at information set \mathcal{S}_1 , then i will bid with a positive probability at s_z . If i does submit her bid $v_\Gamma + d$ at seller s_z , then z remains a high bidder and the standing bid increases to $v_\Gamma + d$. After this event, no trader bids at seller s_z in the continuation because the lowest standing bid never reaches $v_\Gamma + d$.

On the contrary, suppose that in information set \mathcal{S}_1 controlled by buyer i buyer j holds a high bid $b_j > v_\Gamma$ with seller s_j whose standing bid is v_Γ and $b_j \notin B_\Gamma^d$. Then s_j 's standing bid has reached this level before j has submitted her high bid b_j . Therefore, strategy σ^* prescribes that at \mathcal{S}_1 buyer i chooses s_j with zero probability. Since this is true at every information set of G_Γ where the lowest standing bid equals v_Γ and on the continuation path the lowest standing bid never exceeds v_Γ , no buyer will ever bid at s_j after j has submitted her bid b_j . \square

Lemmas A.1 and A.2 imply that if all buyers follow σ^* starting from the initial state Γ_0 in which no buyer has yet submitted a bid, then all trades will take place at price v_m and all buyers with valuations strictly above v_m and all sellers with reserve prices below v_m will trade.

To complete the proof of the theorem, we need to show that no buyer can increase her/his expected payoff in the continuation game G_Γ by deviating from σ^* if all other buyers follow σ^* . This is done below in Lemmas A.3–A.5 which deal with the three possible cases in terms of the effect of the buyer's deviation on the lowest standing bid at the terminal stage: a buyer's deviation causes the lowest standing bid to remain below v_Γ /rise above v_Γ /does not change it, respectively.

Note that a buyer typically does not know the state of the game (in particular, the profile of the other buyers' valuations, their bidding histories, and the array of the high bids) precisely, and hence is not certain about v_Γ and the effect of her deviation on the lowest standing bid at the terminal stage. But the arguments below show that for any v_Γ and any effect of a

buyer's deviation on the lowest standing bid at the terminal stage, a deviation from σ^* does not increase her payoff. Hence, even if a buyer had complete information regarding the state of the game, it would still be optimal for her to follow σ^* , provided that the other buyers do so. Thus, our proof establishes a stronger result that following σ^* constitutes not only a perfect Bayesian equilibrium, but an ex post equilibrium of the bidding game.

The proofs of Lemmas A.3 and A.4 are based on simple supply–demand arguments. Specifically, we show that a buyer can only cause the lowest standing bid at the terminal stage to fall below v_Γ if she herself ends up not trading. Yet, by Lemma A.2, a buyer following σ^* would trade at price v_Γ for sure if her valuation was greater than v_Γ . So a deviation lowering the lowest standing bid is not profitable. Similarly, Lemma A.4 shows that a buyer could cause the lowest standing bid to rise above v_Γ only if she purchases at least as many units as she would purchase had she followed σ^* , and pays higher prices for them.

Finally, Lemma A.5 shows that a deviation from σ^* which does not change the lowest standing bid at the terminal stage cannot be profitable either. This lemma is immediate for any buyer who has not deviated from σ^* by posting multiple high bids or bidding more than one grid point above the standing bid before the path of the bidding game reaches the state Γ . Such buyer gets the best possible outcome in the continuation game G_Γ by following σ^* rather than any alternative strategy that does not change the lowest standing bid at the terminal stage. The proof of the lemma is significantly more complex if a buyer has made one or both described deviations from σ^* prior to state Γ . This buyer may then try to undo the consequences of his earlier deviations by deviating from σ^* in G_Γ also. The extra complexity of the proof is mainly due to the fact that one can no longer rely on simple demand–supply arguments since the lowest standing bid remains the same whether a buyer deviates from σ^* in G_Γ or not. To prove Lemma A.5 we first show that a buyer cannot affect the expected outcome unless she submits new high bids in G_Γ , and then demonstrate that submitting new high bid(s) is unprofitable because it causes both the number of units that she wins and her expected aggregate payment to increase.

Lemma A.3. *Suppose that in G_Γ all buyers other than i follow σ^* , buyer i has valuation v_i and follows some strategy $\sigma \neq \sigma^*$, and the lowest standing bids at the terminal stage is v s.t. $v < v_\Gamma$. Then $v_i \geq v_\Gamma$, in state Γ buyer i does not hold any high bids greater or equal to v_Γ , and at the terminal stage i does not win any units at a price strictly below v_Γ .*

Proof. By hypothesis, the lowest standing bid at the terminal stage does not exceed $v_\Gamma - d$. So, any bidder $j \in D_\Gamma^a(v_\Gamma - d)$, $j \neq i$, must be a winning bidder and on the continuation path j submits bids only with sellers from $S_\Gamma^0(v_\Gamma - d)$. By definition, $\#S_\Gamma^0(v_\Gamma - d) < \#D_\Gamma^a(v_\Gamma - d)$. Therefore, $i \in D_\Gamma^a(v_\Gamma - d)$, so $v_i \geq v_\Gamma$ and i cannot trade with a seller from $S_\Gamma^0(v_\Gamma - d)$. Since all other sellers have high bids or standing bids of at least v_Γ , i cannot trade at a price below v_Γ . \square

Lemma A.4. *Suppose that in G_Γ all buyers other than i follow σ^* , buyer i follows some strategy $\sigma \neq \sigma^*$, and the lowest standing bid at the terminal stage is \hat{v} satisfying $\hat{v} > v_\Gamma$. Then the number of units that i wins and her total payment is at least as large as the maximal*

possible number of units that he would win and the maximal possible total payment that he would make if she follows σ^* .

Proof. First, let us compute the minimal number of units that buyer i wins when she uses σ . Since the lowest standing bid at the terminal stage is \hat{v} and the buyers other than i follow σ^* , a buyer $j \neq i$ holds at most $d_j(\Gamma, \hat{v} - d) + \sum_{v=\hat{v}}^{\bar{d}} g_j(\Gamma, v)$ high bids (recall that $g_j(\Gamma, v)$ is the number of buyer j 's high bids at Γ with sellers whose standing bids are equal to v). To see this, note that $d_j(\Gamma, \hat{v} - d) + \sum_{v=\hat{v}}^{\bar{d}} g_j(\Gamma, v)$ is at least as large as the number of high bids of at least \hat{v} that j holds in state Γ and no less than 1 if $v_j \geq \hat{v}$.

Since $\#(S_\Gamma(v) \setminus S_\Gamma(v - d)) \geq \sum_j g_j(\Gamma, v)$, we conclude that after deviating to σ i wins at least $\#S_\Gamma(\hat{v} - d) - \sum_{j \neq i} d_j(\Gamma, \hat{v} - d) + \sum_{v=\hat{v}}^{\bar{d}} g_i(\Gamma, v)$ units, including at least $\#S_\Gamma(\hat{v} - d) - \sum_{j \neq i} d_j(\Gamma, \hat{v} - d) \geq 0$ units at sellers from $S_\Gamma(\hat{v} - d)$. The last expression is nonnegative because $\#S_\Gamma(\hat{v} - d) \geq AD_\Gamma(\hat{v} - d)$.

Now suppose that buyer i follows σ^* . Then Lemma A.2 implies that the number of units that she wins at sellers from $S_\Gamma(\hat{v} - d)$ is at most $\#S_\Gamma(v_\Gamma) - \sum_{j \neq i} d_j(\Gamma, v_\Gamma) + \sum_{v=v_\Gamma+d}^{\hat{v}-d} g_i(\Gamma, v)$. She also wins $\sum_{v=\hat{v}}^{\bar{d}} g_i(\Gamma, v)$ units at sellers whose standing bids at Γ are at least \hat{v} .

When i follows σ (σ^*), buyers other than i do not place any bids at sellers whose standing bids at Γ strictly exceed \hat{v} (v_Γ). So, to establish the result it is sufficient to show that, when buyer i follows σ , she purchases at least as many units from sellers in $S_\Gamma(\hat{v} - d)$ as when she follows σ^* , i.e.

$$\begin{aligned} \#S_\Gamma(\hat{v} - d) - \sum_{j \neq i} d_j(\Gamma, \hat{v} - d) \\ \geq \#S_\Gamma(v_\Gamma) - \sum_{j \neq i} d_j(\Gamma, v_\Gamma) + \sum_{z=v_\Gamma+d}^{z=\hat{v}-d} g_i(\Gamma, z). \end{aligned} \quad (1)$$

Inequality (1) holds trivially as equality when $v_\Gamma = \hat{v} - d$. Otherwise, note that (1) is implied by the following inequality:

$$\#S_\Gamma(v + d) - \#S_\Gamma(v) \geq g_i(\Gamma, v + d) + \sum_{j \neq i} (d_j(\Gamma, v + d) - d_j(\Gamma, v)) \quad (2)$$

To see that (2) holds, observe the following:

- (i) $\#S_\Gamma(v + d) - \#S_\Gamma(v)$ is the number of sellers whose standing bids in Γ are exactly equal to $v + d$.
- (ii) $g_i(\Gamma, v + d)$ is the number of high bids that buyer i holds in Γ with sellers whose standing bids at Γ are exactly equal to $v + d$.
- (iii) $d_j(\Gamma, v + d) - d_j(\Gamma, v)$ does not exceed the number of high bids that buyer j holds in Γ with sellers whose standing bids at Γ are exactly equal to $v + d$. \square

Lemma A.5.¹⁹ Suppose that in G_Γ all buyers other than i follow strategy σ^* , buyer i follows some strategy $\sigma \neq \sigma^*$, and the outcome is such that the lowest standing bid at the terminal stage is v_Γ . Then i 's payoff is no higher than the payoff that she obtains by following σ^* .

Lemmas A.3–A.5 establish that there is no profitable deviation from σ^* and, thus, complete the proof. \square

Proof of Theorem 2. Consider a seller j with cost c_j . We need to show that seller j 's expected payoff decreases in her reserve price p if $p > c_j$ when the number of traders is sufficiently large. This will be achieved by comparing the expected payoffs that the seller gets when she posts reserve prices equal to p and $p - d$.

By Theorem 1, when all buyers follow strategy σ^* the uniform trading price in the market is equal to v_m , the m th lowest element in \mathbf{v} , the vector of the true buyers' valuations and the sellers' reserve prices. Hence, the uniform trading price is equal to some p^T if and only if the following two conditions hold. First, the number of sellers and buyers whose reserve prices and valuations are strictly below p^T does not exceed $m - 1$. Second, the number of sellers and buyers whose reserve prices and valuations do not exceed p^T is at least m .

Let us fix the strategies of sellers other than j by assuming that all of them set their reserve prices equal to their true costs. Recall the following notation introduced above: $m_1/m_2/m_3$ is the number of buyers whose valuations are strictly below/equal to/strictly above p . Similarly, $n_1/n'_2/n_3$ is the number of sellers, other than seller j , with costs strictly below/equal to/strictly above p . Obviously, $m_1 + m_2 + m_3 = m$ and $n_1 + n'_2 + n_3 = n - 1$. The following two claims hold independently of the profile of strategies used by the sellers other than j .

Claim 1. Suppose that if seller j sets 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 j sets a different reserve price $p' < p_T$. Seller j will trade in both cases.

Proof. By Theorem 1, the trading price is equal to v_m which is not affected by a change in j 's reserve price p as long as $p < v_m$. By Corollary 1, every seller who posts a price below v_m trades. \square

Claim 2. Suppose that if seller j 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 j posts reserve price $p'' > p_T$. Seller j will fail to trade in both cases.

Proof. The trading price is equal to v_m and is not affected by a change in j 's reserve price p as long as $p > v_m$. By Corollary 1, any seller who posts a price above v_m does not trade. \square

¹⁹ The proof of this lemma is available at http://microeconomics.ca/sergei_severinov/proofsinternetauc.pdf.

We will say that price p is *pivotal* when the following condition holds: the trading price is equal to p if seller j sets a reserve price equal to p , all sellers other than j post reserve prices equal to their true costs and all buyers follow strategy σ^* . Thus, p is pivotal if $v_m(\mathbf{b}^m, \mathbf{c}^n_j, p) = p$. Claims 1 and 2 imply that seller j may get a different payoff from setting her reserve price equal to $p - d$ rather than p only if either p or $p - d$ is pivotal.

Let $P(\omega)$ ($P(\omega|\Omega)$) denote the probability of event ω (conditional on event Ω) and $E(y)$ ($E(y|\Omega)$) denote the expectation (conditional expectation given event Ω) of the random variable y . The result of the Theorem follows from the following four lemmas:²⁰

Lemma A.6. *Seller j with cost c_j obtains a higher expected payoff by setting reserve price $p - d$ rather than p , where $p > c_j$, if the following condition holds:*

$$\begin{aligned} &P(p \text{ is pivotal, } p - d \text{ is not pivotal, seller posting } p \text{ fails to trade } |c_j) \\ &\geq P(p \text{ is pivotal, } p - d \text{ is pivotal, seller posting } p \text{ trades } |c_j). \end{aligned} \quad (3)$$

Lemma A.7. *If all buyers follow strategy σ^* and all sellers other than j set reserve prices equal to their true costs, then*

$$\begin{aligned} &P(p \text{ is pivotal, } p - d \text{ is pivotal, seller posting } p \text{ trades } |c_j) \\ &\leq E\left(\frac{1}{n'_2 + 1} \mid m_1 + n_1 = m - 1, m_2 \leq n_1, c_j\right) \\ &\quad \times P(m_1 + n_1 = m - 1, m_2 \leq n_1 |c_j). \end{aligned} \quad (4)$$

Lemma A.8. *If all buyers follow strategy σ^* and all sellers other than j set reserve prices equal to their true costs, then*

$$\begin{aligned} &P(p \text{ is pivotal, } p - d \text{ is not pivotal, seller posting } p \text{ fails to trade } |c_j) \\ &\geq E\left(\max\left\{0, \frac{n'_2 + 1 - m_3 - \min\{0, m_2 - n_1\}}{n'_2 + 1}\right\}\right) \\ &\quad \times |m_1 + n_1 < m - 1 \leq m_1 + n_1 + m_2 + n'_2, c_j) \\ &\quad \times P(m_1 + n_1 < m - 1 \leq m_1 + n_1 + m_2 + n'_2 |c_j). \end{aligned} \quad (5)$$

Lemma A.9. *There exists $N > 0$, independent of p , such that the right-hand side of (4) is less than the right-hand side of (5) if $m = kn \geq N$ (where k is a positive constant).*

References

- [1] S. Anwar, R. McMillan, M. Zheng, Bidding behavior in competing auctions: evidence from eBay, *Europ. Econ. Rev.*, forthcoming.
- [2] L.M. Ausubel, P. Milgrom, Ascending auctions with package bidding, *Frontiers of Theoretical Economics*, BE Press, vol. 1(1), 2002.
- [3] P. Bajari, A. Hortacsu, Winner's curse, reserve prices, and endogenous entry: empirical insights from eBay auctions, *RAND J. Econ.* 34 (2003) 329–355.

²⁰ The proofs of these lemmas are available at http://microeconomics.ca/sergei_severinov/proofsinternetauc.pdf.

- [4] P. Bajari, A. Hortacsu, Economic insights from internet auctions, *J. Econ. Lit.* 42 (2004) 457–486.
- [6] M. Cripps, J. Swinkels, Depth and efficiency of large double auctions, Washington University, St. Louis, 2002.
- [7] E. Dekel, A. Wolinsky, Rationalizable outcomes of large independent private-value first price discrete auctions, *Games Econ. Behav.* 45 (2003) 175–188.
- [8] C. Johnson, On-line auctions will boom through 2007, Forrester Research.
- [9] M. Jackson, J. Swinkels, Existence of equilibrium in single and double private value auctions, *Econometrica* 73 (2005) 93–139.
- [10] V. Krishna, M. Perry, Efficient mechanism design, mimeo, Pennsylvania State University, 2002.
- [11] D. Lucking-Reiley, Auctions on the internet: what is being auctioned and how?, *J. Ind. Econ.* 48 (3) (2000) 227–252.
- [12] P. McAfee, Mechanism design by competing sellers, *Econometrica* 61 (6) (1993) 1281–1312.
- [13] J. Morneau, E-Marketplaces Start to Take Off, TechWeb News, 2001, <http://www.techweb.com>.
- [15] M. Peters, A competitive distribution of auctions, *Rev. Econ. Stud.* 64 (1997) 97–124.
- [16] M. Peters, S. Severinov, Competition among sellers who offer auctions instead of prices, *J. Econ. Theory* 75 (1) (1997) 141–179.
- [17] A. Roth, A. Ockenfels, Last minute bidding and the rules for ending second price auctions: theory and evidence from a natural experiment on the internet, NBER Discussion paper, 2002.
- [18] A.E. Roth, M.A. Sotomayor, *Two-Sided Matching—A Study in Game, Theoretic Modelling and Analysis*, Econometric Society Monographs, Cambridge University Press, Cambridge, MA, 1990.
- [19] A. Rustichini, M.A. Satterthwaite, S.R. Williams, Convergence to efficiency in a simple market with incomplete information, *Econometrica* 62 (5) (1994) 1041–1063.
- [20] M.A. Satterthwaite, S. Williams, The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large, *Review of Economic Studies* 56 (1989) 477–498.
- [21] R. Wilson, Incentive efficiency of double auctions, *Econometrica* 53 (1985) 1101–1115.
- [22] T. Wilson, T. Mullen, E-business exchanges fight for survival, *InternetWeek*, 2000, <http://www.internetwk.com>.
- [23] S. Williams, A characterization of efficient, Bayesian incentive compatible mechanisms, *Econ. Theory* 14 (1999) 155–180.