Shill bidding and optimal auctions.*

Sergei Izmalkov†

September 30, 2004

Abstract

A single object is for sale to N asymmetric buyers in the independent private values setting by means of an English auction. We consider implications of seller's active participation in the bidding process, or shill bidding. The main result is that there exists an equilibrium of the English auction with shill bidding that is outcome equivalent to the optimal mechanism of Myerson (1981). We also show that common knowledge requirements for existence of the optimal equilibrium can be significantly reduced.

1 Introduction

How to sell an object and profit the most? A famous answer is given by Myerson (1981) for the asymmetric independent private values (AIPV) setting: buyers' valuations of the object are private (only known to them), do not depend on what the others know, and are randomly independently drawn from not necessarily identical distributions. In a direct mechanism the buyers report to the seller their values and the allocation and the payment rules are as follows. The seller calculates virtual valuations for each of the buyers and sells the object to the buyer with the highest positive virtual valuation.

*Very pleminінary. Suggestions and comments are very welcome. The most recent version of the paper is available at http://econ-www.mit.edu/faculty/index.htm?prof_id=izmalkov&type=paper.

†Address: 50 Memorial Drive, Department of Economics, MIT, E52-252C, Cambridge, MA 02142. E-mail: izmalkov@mit.edu.

If all virtual valuations are negative the object is not sold.[1] When sold, the price the winner pays is the lowest possible value she could have had and still win.

In the symmetric setting, when buyers' valuations come from identical distributions, the optimal mechanism is extremely simple: sell to the buyer with the highest value as long as it exceeds a certain minimal level, a reserve price, at a price that equals the higher of the second highest value and the reserve. The optimal mechanism can be implemented by a variety of simple selling procedures, by the first- or second-price sealed bid auctions, English or Dutch auctions with appropriately chosen reserve price, and many others (Riley and Samuelson (1981)).[2]

When buyers are asymmetric, the optimal mechanism is neither neutral no anonymous, it depends on characteristics of the object and of the buyers and it treats buyers differently.[3] This fact together with the absence of alike selling mechanisms in practice is the main source of dissatisfaction with the obtained characterization and is a point of departure of several studies that either impose further restrictions on the solution concept (strengthening Bayesian incentive compatibility (IC) to dominant strategy IC, Chung and Ely (2003)), or on the seller (limited commitment power, Vartiainen (2002), Skreta (2003)), or consider a limited family of mechanisms to choose from (simple sequential auctions, Lopomo (1998)).

In this paper we show that the optimal mechanism can be implemented by

[1] Virtual valuations can be thought of as marginal revenues a monopolist (the seller) obtains from selling to a buyer of a given value (see Bulow and Roberts (1989)). For formal definitions and a derivation of the optimal mechanism see Section 2.

[2] We employ a weak notion of implementation, requiring only existence of an equilibrium that has desired allocation and payment rules.

[3] Wilson (1987) writes, "The optimal trading rule for a direct revelation game is specialized to a particular environment. For example, the rule typically depends on the agents' probability assessments about each other's private information. Changing the environment requires changing the trading rule. If left in this form, therefore, the theory is mute on one of the most basic problems challenging the theory. I refer to the problem of explaining the prevalence of a few simple trading rules in most of the commerce conducted via organized exchanges. A short list—including auctions, double auctions, bid-ask markets, and specialists trading—accounts for most organized exchange. ... The rules of these markets are not changed daily as the environment changes; rather they persist as stable, viable institutions. As a believer that practice advances before theory, and the task of theory is to explain how it is that practitioners are (usually) right, I see a plausible conjecture: These institutions survive because they employ trading rules that are efficient [Pareto] for a wide class of environments."

the most common of auction formats, the open ascending price, or English auction. The key is to allow the seller *shill bid*—freely participate in the bidding process.

Usually, when auctions are modelled the seller is either in possession of full powers of mechanism designer or is completely ignored. The corresponding games and their solution concepts are defined for buyers-bidders only. Even when allowed to act, the boundaries of the seller's actions are explicitly defined—affecting bidders' information or setting a reserve price are typical examples—and are separated in time with actions of the buyers, so that the game can be analyzed sequentially.

Suppose that the seller's choice of trading rules is restricted to a particular auction form. Given that the seller's objective is to maximize revenue, it seems only appropriate to grant her full freedom to act within the rules of the auction.[4] The model that restricts seller's activities may *a priori* be inadequate in its revenue properties. Unlike sealed-bid formats, where the seller's bid is essentially a reserve price, an English auction permits the seller to adjust her behavior depending on the observed bidding.[5]

Our main result is that the English auction with shill bidding has an equilibrium that is outcome equivalent to the optimal mechanism. We also show that if the seller's valuation is also private, the seller would like to set an open reserve price in addition to possible shill bidding later. Thus, both open reserve prices and shill bids coexist and serve different purposes. The open reserve price is intended to reveal the seller's valuation, which may explain why in practice reserve prices are sometimes associated with reservation values. Shill bids, or secret reserve prices, are set to extract some of the buyers' surplus.[6]

[4]Even if the rules of a given auction do explicitly forbid seller's participation it is still instructive to know whether the seller has an incentive to participate or manipulate the rules.

[5]In essence, the English auction is characterized by an open bidding process where the price is constantly and openly raised until the object is sold. Exact manner in which this is done varies: in some auctions the bidders cry out their bids, and the auction ends when no one is willing to offer more than the standing bid; in others the auctioneer calls out the price and raises it in small increments as long as at least two bidders are willing to pay the current price. We use a price-clock model of the English auction originated in Milgrom and Weber (1982), for details see Section 3.

[6]A shill bid is effectively an unobserved minimal price at which the seller is willing to sell, or a secret reserve price. In literature and in practice, a secret reserve price often has a more narrow meaning, it may be required to be set in advance, and sometimes its

The *optimal* equilibrium is not unique, and in it the buyers play weakly-dominated strategies—they exit before the price reaches their values whenever there are more than two active participants in the auction. Another equilibrium, where the buyers follow a weakly dominant strategy—remain active until the price reaches own value—exists as well. In both the optimal and the *dominant* equilibria the seller stays active a long as two other buyers are still bidding and once one of them exits, she either exits immediately or shill bids: remains active until a certain price based on identity of the remaining buyer.

The selection toward an equilibrium in which (some) players follow their dominant strategies seems only logical. And, while indeed the dominant equilibrium seems to be the most plausible solution, we provide two arguments against a selection on grounds of dominance for the game at hand. Both arguments directly or indirectly account for a fact that the seller is the (only) player able to commit.[7] By the very notion of dominance, for a buyer, to check whether one strategy dominates another she needs to compare her payoffs against any possible profile of strategies of the rest of the players, including the seller. The seller, however, once committed, has only one possible strategy. If we look at dominance after commitment, then once the seller's strategy is fixed, buyers' strategies in both equilibria are undominated.[8]

The second argument recognizes that the choice and the properties of a particular solution to an economic interaction is fundamentally linked to the way the interaction is modeled as a game. While the robustness of a particular equilibrium to model misspecifications is a separate equilibrium refinement concept, we show how it may affect the strength of selection on grounds of dominance. We do so by considering a possibility of reentry and by analyzing implications of post-auction negotiations if the seller retains the object. In the first case, dominant strategies no longer exist and both considered equilibria are undominated. In the second case, the dominant equilibrium ceases to exist.

A particularly attractive property of the dominant strategies equilibrium and also of an ex post equilibrium concepts is that they allow for a signifi-

existence has to be announced as well. See Katkar and Lucking-Reiley (2001) for eBay rules on secret reserve prices.

[7]This assumption is central in deriving the optimal mechanism, we require it as well.

[8]In the dominant equilibrium, a strategy to stay until the price reaches own value is no longer weakly dominant for buyers' with negative virtual valuations. A strategy "exit at $p = 0$" provides the same payoff no matter what strategies the other buyers use.

cant reduction in common knowledge requirements imposed on the players.[9] These equilibrium concepts allow a player to be completely ignorant of how the values of all the players are distributed and even of whether these distributions come from the common prior (Bergemann and Morris (2003)), and so the common knowledge assumption can be significantly reduced. In the English auction with shill bidding there are no dominant strategies or ex post equilibria. Recognizing a special role of the seller, we propose two equilibrium concepts that are asymmetric: *a buyers' dominant strategies equilibrium* and *a buyers' ex post equilibrium*. The introduced difference is that the corresponding equilibrium requirements on buyer's dominant strategies equilibria. Exact minimal common knowledge requirements for their existence can be pinned down. In the dominant equilibrium, a buyer needs to know only her own value, the seller has to know all the distributions. In the optimal equilibrium, a buyer needs not to know other buyers' (and the sellers') value distributions, her own distribution has to be a common knowledge between her and the seller.

Auctions with shill bidding and, in particular, revenue effects of shill bidding in English auctions have been analyzed before. The closest papers are Graham, Marshall, and Richard (1990) and Graham, Marshall, and Richard (1996). They study English auctions where the seller's reserve price can be a function of the highest observed bid. These papers are the first to recognize that the optimal shill bid may be a function of the history when the bidders are heterogeneous and to show that shill bidding can enhance sellers' revenue. Their main emphasis is on modelling uncertainty a seller might have about identities of the bidders and its effects on the seller behavior. Wang, Hidvegi, and Whinston (2001) consider a symmetric setup with non-monotone virtual valuations and with uncertain number of bidders. The optimal reserve price in this case depends on the number of bidders. They show that the English auction obtains the highest possible revenue since effectively the seller can observe the actual number of bidders and set her reserve price—shill bid—as a function of that number. Several papers study shill bidding in common values settings.

[9] An equilibrium is in dominant strategies if for each player i there exists a strategy s_i such that for any other strategy s_i' and any profile of strategies of the other players, $\mathbf{s}_{-i}$, $u_i(s_i, \mathbf{s}_{-i}) \geq u_i(s_i', \mathbf{s}_{-i})$. A Bayesian-Nash equilibrium is ex post if it remains an equilibrium even if realized types $(x_1, x_2, ..., x_n)$ become common knowledge, $u_i(s_i(x_i), \mathbf{s}_{-i}(\mathbf{x}_{-i})) \geq u_i(s_i', \mathbf{s}_{-i}(\mathbf{x}_{-i}))$ (compare with B-N equilibrium condition, $E_{\mathbf{x}_{-i}} u_i(s_i(x_i), \mathbf{s}_{-i}(\mathbf{x}_{-i})) \geq E_{\mathbf{x}_{-i}} u_i(s_i', \mathbf{s}_{-i}(\mathbf{x}_{-i})))$.

A recent paper by Caillaud and Robert (2003) presents another detail-free mechanism implementing the optimal auction. Their idea is to make a potential winner to offer a price and let other buyers challenge the price at will—effectively monitoring each other. If there is no challenge the object is sold, if there is a challenge the winner pays a penalty and the object is sold at a new price if she accepts the challenge or not sold at all. A potential winner is determined by means of an English auction, where bidders compete for the right to make an offer to the seller and their exiting prices serve as a revelation device for their values. In an equilibrium buyers bid according to their virtual valuations, and later in the challenge stage, the inferred values serve as levels for offered prices.[10]

The rest of the paper is organized as follows. Section 2 introduces the considered setup (AIPV setting) and Myerson's optimal mechanism. Section 3 describes the English auction with shill bidding and presents two equilibria of it, one of which is the optimal. Section 4 discusses several extensions of the analysis, and, in particular, looks whether a requirement of common knowledge of buyers' distributions can be relaxed.

2 The optimal mechanism

2.1 The setup (AIPV setting)

This is a typical setup of many independent private values auction models. A seller has a single object to sell to N buyers. Each buyer's (expected) valuation for the object is fully determined by her private information. All buyers are risk neutral and have quasi-linear utility functions. Without loss of generality we can suppose that buyer i's type, or private signal, $x_i \in [0, w_i]$, is equal to her value. The signals, $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ of all the buyers are drawn from independent but not necessarily identical distributions, $x_i \sim F_i$, with density $f_i > 0$ on $[0, w_i]$. The value of the seller, x_0, is assumed to be equal to 0 (later we will relax this assumption). For now assume that all of the above is commonly known among the seller and the bidders.

[10] Caillaud and Robert (2003) provide a few other mechanisms that implement the optimal mechanism in AIPV setting. One is trough ascending-bid auction followed by a special payment rule. There the bids would coincide with virtual valuations, but the payments are identity based. A second mechanism implements the optimal auction with individual price clocks that have different speeds.

2.2 The optimal auction

A characterization of the auction that maximizes seller's expected revenue—one of the classic results of the auction theory—is provided by Myerson (1981).[11] By revelation principle the analysis can be restricted to direct mechanisms. A direct mechanism asks each bidder to report her value and consists of a pair $(\mathbf{Q}, \mathbf{M})$: an allocation rule $\mathbf{Q} : \mathbf{x} \to [0,1]^N$, where $Q_i(\mathbf{x})$ is the probability that buyer i receives the object given reports $\mathbf{x}$; and a payment rule $\mathbf{M} : \mathbf{x} \to \mathbb{R}^N$, where $M_i(\mathbf{x})$ is the expected payment by buyer i.

A virtual valuation of buyer i is

$$\psi_i(x_i) \equiv x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}. \tag{1}$$

The analysis below is presented for the main (regular) case, when all the virtual valuation functions are increasing. If they are not, the virtual valuations need to be "ironed out" first to obtain non-decreasing quasi-virtual valuations $\psi_i^*(x_i)$, see Myerson (1981) for the description of the procedure. Then, the optimal mechanism is based on the quasi-virtual valuations, the results of the paper straightforwardly extend.[12]

The optimal auction $(\mathbf{Q}^*, \mathbf{M}^*)$ specifies that the object is allocated to the buyer with the highest marginal valuation provided it exceeds 0. Only the winner pays, and the price is equal to the lowest value the winner can have and still win. Formally,

$$Q_i^*(\mathbf{x}) = \begin{cases} 1, \text{ if } \psi_i(x_i) > \max_{j \neq i} \psi_j(x_j) \text{ and } \psi_i(x_i) \geq 0; \\ 0, \text{ if } \psi_i(x_i) < \max\{0, \max_{j \neq i} \psi_j(x_j)\}. \end{cases} \tag{2}$$

$$M_i^*(\mathbf{x}) = \begin{cases} \psi_i^{-1}(\max\{0, \max_{j \neq i} \psi_j(x_j)\}), \text{ if } Q_i^*(\mathbf{x}) = 1; \\ 0, \text{ if } Q_i^*(\mathbf{x}) = 0. \end{cases} \tag{3}$$

In the case when two or more buyers share the highest virtual valuation and it is non-negative the object is assigned to either of them at random and the winner pays her value.[13]

---

[11] Our presentation follows one in Krishna (2002).

[12] The only technical isssue is to ensure that all the types of a given buyer with identical quasi-linear valuations have the same probabilities of winning. In an English auction, in the considered equilibria, this is achieved automatically since all the types in question are bidding in the same way.

[13] Optimal mechanism can be derived as follows. Given $(\mathbf{Q}, \mathbf{M})$, let $U_i(z, x_i) = q_i(z)x_i -$

3 English auctions with shill bidding

3.1 The English auction

Since Milgrom and Weber (1982), an open ascending price, or English, auction has been modeled as a price-clock auction. The price is shown on the clock, it starts at 0 and continuously increases. As the price increases bidders indicate whether they are still willing to purchase the object at the current price, are active, or not. A bidder can stop bidding, or exit the auction, at any price, but once this happens, she cannot become active again. Exit prices and identities of the bidders are commonly observed. The clock stops when only one bidder remains active. She wins the object and pays the price shown on the clock. If all remaining bidders exit simultaneously the winner is selected randomly among them.

A strategy of bidder i of type x_i can be described as a collection of intended exit prices: when to drop out if no other active bidder exits first, given that i is active at p, and given current public history of bidding—a collection of prices at which other bidders have dropped out. To avoid excessive notation we would simply use $p_i(x_i)$ as an intended exit price and p_i as an actual exit price. The public history, if relevant, would be clear from the context.[14]

The original price-clock model does not allow the clock to be stopped

---

$m_i(z)$ be expected payoff of player i with true signal x_i and report z; here $q_i(z_i)$ and $m_i(z_i)$ are the probability of winning and the expected payment of player i reporting z. Then $u_i(x_i) = U_i(x_i, x_i) = \max_z\{q_i(z)x_i - m_i(z)\}$ by IC, and so $u_i(x)$ is convex and differentiable almost everywhere. IC implies that for any y and x_i, $u_i(y) \geq u_i(x_i) + q_i(x_i)(y - x_i)$. Thus, if u is differentiable at x_i, then $u'(x_i) = q_i(x_i)$. Convexity implies monotonicity of $q_i(x)$. We can write $u_i(x_i) = u_i(0) + \int_0^{x_i} q_i(y_i)dy_i$, and so $m_i(x_i) = q_i(x_i)x_i + m_i(0) - \int_0^{x_i} q_i(y_i)dy_i$.

The expected payment of player i to the seller is $E_{x_i}[m_i(x_i)] = m_i(0) + \int_0^{w_i} x_i q_i(x_i) f_i(x_i)dx_i - \int_0^{w_i}\int_0^{x_i} q_i(y_i) f_i(x_i) dy_i dx_i$. After changing the order of integration and reducing, the last term is equal to $\int_0^{w_i} (1 - F_i(y_i))\, q_i(y_i) dy_i$. Combining, we obtain $E_{x_i}[m_i(x_i)] = m_i(0) + \int_0^{w_i} \left(x_i - \frac{1-F_i(x_i)}{f_i(x_i)}\right) q_i(x_i) f_i(x_i) dx_i$. Summing up, and remembering what $q_i(x_i)$ is, the expected revenue to the seller is

$$E[R] = \sum_{i=1}^{n} m_i(0) + \sum_{i=1}^{n} \int_{\mathcal{X}} \left(x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}\right) Q_i(\mathbf{x}) f(\mathbf{x}) d\mathbf{x}.$$

Optimal allocation and payment rules immediately follow. Monotonicity of virtual valuations guarantees monotonicity of $q_i(x)$.

[14] For a detailed formal treatment of the game the reader is referred to Krishna (2003).

during the auction. This precludes strategies "exit as soon as possible" to be played after observing some other bidder dropping out, simply because such strategies are not well-defined. In our analysis, such a strategy may be a best reply. To ensure that it is well-defined we modify the model by allowing the clock to be stopped after an exit. Without stopping the clock, for example, the dominant equilibrium presented below may not exist (see Footnote 18).

What exactly happens once the clock is stopped is not crucial as long as the following two principles hold: a bidder who initiates the clock stop (exits first) cannot be the winner unless all the remaining bidders exit simultaneously; all the bidders are treated the same way. We consider the following procedure. The clock stops when one or more bidders exit. If multiple bidders exit simultaneously, only one of them, chosen randomly, becomes inactive, others remain active. Once the clock is stopped, all remaining active bidders simultaneously choose whether they would like to remain active or exit. If no bidders exit and at least two are active, then the clock is restarted. Otherwise, if one or more bidders exit, then the procedure is repeated, one bidder is randomly selected to drop out, and so on. The auction ends when only one bidder remains active. Note that although we let the bidders exit only one-by-one, if all the bidders decide to exit at the same price, each of them has the same probability of winning the auction—the same way "ties" are resolved in the original model.

To complete our specification of the game we need to describe in greater detail how the auction starts. Each bidder decides whether to enter the auction or not. The clock is initially stopped at 0. All the bidders who enter the auction are considered active. To start the clock the same procedure as if a bidder exited at 0 is conducted.[15]

3.2 Shill bidding

Bid padding, phantom bidding, bidding of the wall, lift-lining, trotting, running, setting hidden reserve prices—these are examples of the seller's or auc-

[15] To account for the modifications made to the model we can set $p_j = -1$ for any bidder j that does not enter the auction. Also, both the intended and actual exit prices can be equal to the current price. Strictly speaking, public histories can no longer be summirized by the sequences of exits. Bidder i deciding to exit only after bidder j exits and stops the clock, and exiting simultaneously (but after) with j are two different histories, provided all attempts to exit are observable. This, however, is irrelevant for the analysis that follows, so we would still associate public histories with sequences of exits.

tioneer's activities that involve active participation in the bidding process. For an excellent account of these practices see Cassady (1967). Shill bidding is a modern composite term that includes these activities and much more. A leading on-line auction market, eBay, defines shill bidding as the deliberate placing of bids to artificially drive up the price of an item. For a more detailed account of shill bidding practices see Appendix A.1.

We define shill bidding as an unrestricted seller's participation in the bidding process within the rules of the game. More precisely, the seller is allowed to act as any other bidder during the auction. The sellers' identity is observed during the auction.[16] If the seller ends up winning the object it is assumed that she remains in possession of it and no further activity is possible—the object goes off the market. Whenever the seller's strategy prescribes her to remain active with only one other bidder active, and so the seller is risking retaining the object, we would say that the seller is shill bidding. We associate the seller with bidder 0, the set of all buyers is denoted as $\mathcal{N}$.

3.3 Two equilibria

Even without shill bidding, an English auction admits many Bayesian-Nash equilibria. Two examples are an asymmetric equilibrium, where one of the buyers remains active forever, while the others exit immediately at $p = 0$ (or do not enter at all), and the dominant strategies equilibrium, where every buyer follows "remains active until the price reaches own value" strategy. The latter equilibrium, being in dominant strategies, is usually selected, and all comparative analysis, e.g. studying revenue and efficiency properties, is conducted based on this equilibrium.

The English auction with shill bidding has multiple equilibria as well. It is not our goal to describe the full set of equilibria,[17] which itself depends on the details of a particular specification of the game. Instead we focus on revenue implications of allowing shill bidding. Here we present two equilibria that stand out among the rest: one is an extension of the dominant strategy equilibrium of the ordinary English auction, the other provides the seller with

[16] We assume that each player, a buyer or the seller, can send only one representative to the auction (or bid herself), and that identities of the players are readily identifiable. Later we will discuss implications of relaxing this assumption.

[17] To the best of our knowledge this is yet to be done for the English auction without shill-bidding.

the same revenue as the optimal mechanism.

Proposition 1 *There exists an equilibrium of the English auction with shill bidding in which each of the buyers remains active until price reaches her value, $p_j(x_j) = x_j$. The seller's strategy calls on her to remain active as long as at least two other bidders are active. Once only one other bidder remains active, say buyer i at price p_{-i}, the seller decision is based on a personal reserve price for buyer i, $\psi_i^{-1}(0)$. The seller remains active until the price reaches $\psi_i^{-1}(0)$, or exits immediately if the current price is higher, $p_0 = \max\{p_{-i}, \psi_i^{-1}(0)\}$.*

Proof. To stay until the price reaches the value is a weakly dominant strategy for each buyer. Given that the probability of the bidders exiting simultaneously and leaving the seller with the object is equal to 0, it is optimal for the seller to stay as long as at least two other bidders are active. Then, once only one bidder, say i, remains active at p_{-i}, the seller can either exit immediately and guarantee a sale at $p = p_{-i}$ or stay longer with the hope of getting a higher price at a risk of no sale. If the seller decides to stay until $p \leq w_i$, the seller's expected revenue is

$$ER = \Pr(x_i > p \mid x_i > p_{-i})p_0 = \frac{1 - F_i(p)}{1 - F_i(p_{-i})}p. \tag{4}$$

It is maximized (unconstrained) at $p^* = \psi_i^{-1}(0)$. Therefore, if $p^* \leq p_{-i}$, the seller is better off exiting immediately, otherwise she should stay until the price reaches the value of bidder i that corresponds to zero virtual valuation. Note that, given the seller's strategy, bidder i cannot gain by exiting immediately at p_{-i}.

Alternatively, the strategy of the seller can be defined as stay until $p_0 = \max\{\max_{i \in \mathcal{A}} \psi_i^{-1}(0), p\}$, where $\mathcal{A}$ is the set of *buyers* active at p.[18] ■

We would refer to this equilibrium as *the dominant* equilibrium. Note that the seller submits a shill bid whenever the imputed virtual valuation of the remaining buyer is below 0. If submitted, a shill bid depends on the identity of the remaining buyer but not on the price at which the last exit have happened.

[18] Suppose that $\mathcal{A}$ consists of only two buyers, i and j, and that $\psi_j^{-1}(0) < p < \psi_i^{-1}(0)$. Each of the buyers can be the first to exit. If buyer j exits first, the seller would remain active, if buyer i does, the seller would like to exit as well. If the clock is not stopped such strategy is not feasible.

Remark 2 *There is some freedom in specifying the seller's strategy in the dominant equilibrium. For example, if the seller knows the current price is high enough so that she will not shill bid against any of the remaining buyers, she can as well exit now. Still, the seller must remain active as long as at least one of the remaining buyers has the personal reserve higher than the current price. Consider a situation with two buyers active, say* i *and* j, *and* $\psi_j^{-1}(0) < p < \psi_i^{-1}(0)$, *so that only buyer* i*'s personal reserve exceeds the current price. Then, if buyer* i *exits, the seller wants to exit immediately. If the clock is not allowed to stop such a strategy is not well-defined, and so, the dominant equilibrium ceases to exist. It does not seem appropriate to loose a reasonable equilibrium on a technicality.*

Proposition 3 *There exists an equilibrium of the English auction with shill bidding with the same allocation and payment rules as the optimal mechanism.*

We would refer to this equilibrium as *the optimal* equilibrium.

Proof. First, we describe the equilibrium construction. The strategies of the buyers are defined as follows. Each buyer j's strategy is: enter the auction only if $\psi_j(x_j) \geq 0$; as long as more than two bidders including j are active, stay until the price reaches virtual valuation, $p_j(x_j) = \psi_j(x_j)$; when only one other bidder is active, stay until the price reaches true valuation, $p_j(x_j) = x_j$. The seller's strategy is: enter the auction; remain active until only one buyer remains active; stay active (shill bid) until $p_0 = \psi_i^{-1}(p_{-i})$, where $i \in \arg\max_{j \in \mathcal{N}} \psi_j(x_j)$ is the last remaining buyer, which happens at $p_{-i} = \max\{\max_{j \neq i} \psi_i(x_i), 0\}$. Note that if no buyer enters the auction, the seller is the only active player at $p = 0$, she wins and pays 0.

It is straightforward to verify that the allocation and payments resulting from these strategies coincide with the optimal allocation and payment rules. To verify that this is an equilibrium, observe that no buyer can increase her payoff. Indeed, only the winner pays. To win a buyer, say j, must end up as one of the two final bidders together with the seller. If the buyer does not have the highest virtual valuation that exceeds 0, the seller would stay active until the price—an inverse of $p_{-j} = \max\{\max_{j \neq i} \psi_i(x_i), 0\}$—that exceeds that buyer's true value, and so it is not profitable for the buyer to win. If the buyer does have the highest virtual valuation that exceeds 0, she pays the lowest possible price given the strategies of the seller and of the other buyers.

A follows from the solution to (4), if $p_{-i} = 0$, it is best for the seller to stay until $\psi_i^{-1}(0)$. At $p^*_{-i} > 0$ the seller infers that the value of buyer i is at least $\psi_i^{-1}(p_{-i})$, so it is "safe" to stay until then, but given that $\psi_i(\psi_i^{-1}(p_{-i})) = p_{-i} > 0$, the optimal strategy for the seller is to exit at $p_0 = \psi_i^{-1}(p_{-i})$. ■

Thus, the ordinary English auction is a detail-free mechanism that implements the optimal auction. As a trading rule it does not depend neither on the characteristics of the object nor on the characteristics of the buyers.

3.3.1 Reserve prices

Here we show that if the seller's value is not known a priori and is a private information to the seller, then the optimal auction can still be implemented by the English auction with shill bidding where the seller uses an open reserve price.

Suppose that the seller's value x_0 is her private information and is distributed as $x_0 \sim F_0[0, w_0]$. The virtual valuations (1) need to be redefined,

$$\psi_i(x_i; x_0) \equiv x_i - \frac{1 - F_i(x_i)}{f_i(x_i)} - x_0. \tag{5}$$

They now reflect the fact that in the case the object is not sold, the value to the seller is x_0. The optimal mechanism consists of allocation (2) and payment (3) rules based on the redefined virtual valuations. Note that if we consider a general mechanism design problem where the seller is not a mechanism designer (whose goal is still to generate the highest expected revenue for any given type of the seller) and have to report (in a direct mechanism) her true value, the incentive compatibility constraints of the seller are trivially satisfied. Given the type of the seller, that type obtains the highest possible expected payoff, and so has no incentive to lie.

Proposition 4 *There exists an optimal equilibrium of the English auction with shill bidding and reserve price. In it the reserve price is set at the reservation value of the seller, and the resulting allocation and payment rule coincide with the optimal allocation and payment rules.*

Proof. By setting reserve price at her reservation value the seller effectively makes it a common knowledge among the buyers. The result trivially follows. Note that the seller, since she obtains her first best in terms of expected revenue, cannot profitably deviate. ■

Thus, we have found that hidden reserve prices and open reserve prices are not only different in terms of the decisions made at the moment they are set (based on informational differences involved), but can also serve simultaneously as two instruments available to the seller and used with completely different purposes: to extract more revenue and to reveal the seller's private information. This result might shed some light on why the bidders tend to associate the reserve price with the seller's reservation value in the real-life English auctions.

4 Discussion

4.1 Commitment and dominance

One strategy dominates another if it guarantees a higher payoff against any possible play of the opponents. This is the essence of the notion of dominance. In a formal definition, 'any possible play' and 'a higher payoff' have to be specified. For example, weak dominance requires at least as high payoff against any possible collection of strategies and a strictly higher payoff against at least one set of strategies of the rest of the players.

In the English auction with shill bidding each buyer has a weakly dominant strategy—to remain active until the price reaches own value. So, in particular, buyers' strategies in the optimal equilibrium are weakly dominated. Whenever there are more than two active bidders, buyer i of type x_i cannot loose if instead of exiting at the price that equals her virtual valuation she remains active until $p_i = x_i$. She gains if all the other bidders, including the seller, exit simultaneously at some $p < x_i$.

In the above definition of weak dominance and in its application, 'any possible play' is translated as any possible collection of strategies of the rest of the players. This, however, seems to be at odds with the assumption that the seller is able to commit. Naturally, once the seller commits to a given strategy, 'any possible play' by the seller consists of this strategy only. If this version is used in the definition of weak dominance, then buyers' strategies in the optimal equilibrium become undominated. They provide the same payoff to the buyers as "stay until the price reaches own value" strategies. Indeed, to win a buyer needs to outbid the seller. Given the seller's strategy, if the price reaches the buyer's virtual valuation while some other buyers are still active, then the seller, once left one-on-one with the buyer, will remain active at least until the price reaches the true value of the buyer. Thus, if wins, the buyer has to pay at least her value.

Similarly, in the dominant equilibrium, the buyers' strategies are no longer weakly dominant. More precisely, buyers with virtual valuations below 0 cannot win and obtain positive profits due to the seller's shill bidding. For them, the strategy "stay until the price reaches own value" is only undominated, it is equivalent to the strategy "exit at $p = 0$."

4.2 Model misspecifications

Defect is the dominant strategy in the Prisoners' Dilemma (PD) game. Does that mean that prediction (Defect, Defect) is the most plausible for an economic conflict that is modeled by PD? The answer is *Yes* if we believe that PD is the correct model, but in general it would depend on the conflict at hand. The joint defection outcome would be less likely to occur if there were features of the conflict situation that undermine dominance properties of the Defect strategy. For example, this may happen if the parties can jointly commit to *ND (C)* strategy, or they have other options (bribing, partial defection, side payments, ...), or the game is the part of a repeated interaction.[19]

In an exact parallel, we may ask how unmodeled features of real-life English auctions may affect plausibility of equilibria based on dominant strategies. The very existence of dominant strategies is tied up to the irrevocable exits assumption of the Milgrom & Weber model. This assumption ensures that every player has only one action in the game—exit the auction. So, in particular, when only two bidders are active, neither of them has an opportunity to react to an action of the opponent—the auction ends once only one bidder remains active. In most real-life ascending auctions stopping rules are much more flexible. Typically, any bidder can raise the would-be-final price if she wishes to do so. This provides ample opportunities for retaliation and effectively prevents existence of dominant strategies. For example, if we consider an English auction with reentry model from Izmalkov (2003), then both the dominant and the optimal equilibria still exist, and both are in undominated strategies. the strategy "stay until the price reaches own value" is no longer dominant.[20]

[19] For more on Prisoners' Dilemma and, in particular, on how cooperation can emerge, see Axelrod (1984).

[20] To complete both equilibriua it suffices to treat any instance of reentry as if a bidder that becomes active again had never dropped out. Then neither the buyers nor the seller

Another assumption that is crucial for existence of dominant strategies is the no post auction activity assumption—no resale or any additional transfers can happen after the auction ends. By allowing resale, the incentives within a game change quite significantly in addition to creation of new opportunities to act. A buyer, for example, might consider exiting immediately and then making a take-it-or-leave-it offer to the winner. The seller, provided all the buyers' play as if there were no resale, may choose to remain active forever with the hope of revealing the true valuations of all the buyers and then selling the object again to the buyer with the highest valuation and extracting the full surplus. At the same time, realizing that the buyers may adjust, the seller may choose to commit not to resale the object.

Here, we present an example of how the post auction activity can be introduced into the model so that it changes significantly the equilibrium properties of the game. Suppose that a resale is still not possible, but if the object is not sold, there is a chance that the seller can make a take-it-or-leave-it-offer at the end of the auction to any of the auction's participants. How exactly this is modeled is not important as long as the probability of a buyer receiving an offer from the seller is positive whenever the deal is profitable to the seller.

Then, the dominant equilibrium no longer exists. Suppose that the realized valuations of the buyers are all different and are such that the seller shill bids and wins the object. Consider the buyer, say j, whose exit leaves only the seller and one other buyer active. The last remaining buyer, say i, must have a negative virtual valuation for the seller to win. Buyer j, then, when exiting, knows that the probability of i having a negative valuation is positive, and so a probability of no sale is positive, and so a probability of receiving an offer from the seller is positive. Given that, buyer j would like to shade her bid, so that it does not reveal her true valuation. It is no longer optimal for her to remain active until price reaches her valuation. In

would reenter in either equilibrium during the auction. The buyers' strategies from the optimal equilibrium are no longer dominated simply because any buyer can come back if there is a profit to it. Moreover, the strategy "stay until the price reaches own value," $p_i = x_i$, is no longer "no worse" than any other strategy. For example, suppose all the other bidders follow the following strategy: exit at $z + \varepsilon$, if i have exited at z; exit at w_i, otherwise. Then, for sufficiently small $\varepsilon > 0$ against this profile of strategies, for buyer i of type $x_i > z + \varepsilon$ it is worse to stay until $p_i = x_i$ than to exit at $p_i = z$ and reenter at $p_i = z + \varepsilon$. The argument still holds even if we fix the strategy of the seller from one of the equilibria.

general, a buyer, if considers a probability of receiving a take-it-or-leave-it offer by the seller to be positive will not follow a strategy that reveals her valuation. Thus, in any monotone equilibrium—such that valuations of the active buyers who do not win are revealed—the object needs to be sold with probability one whenever there are active buyers. Indeed, if the object is not sold with probability 1 while there are active buyers, the seller must be shill bidding and retaining the object on some occasions. An (optimal) shill bid by the seller is essentially a take-it-or-leave-it offer to the last remaining buyer. If the object is not sold, that is the offer is not accepted, the seller will not be making another offer to the the same buyer. As a result, at least one other active buyer if any, will be receiving an offer, so her strategy must not be revealing. Note that, the optimal equilibrium still exists even if post-auction offers are possible. Although it is in revealing strategies, the seller always sells an object whenever there are active buyers.

4.3 Wilson's doctrine

The emphasis on design of detail-free mechanisms, operating in a variety of scenarios and independent of particular details of a scenario at hand, is the first part of the Wilson's doctrine. In an auction setting, for example, a detail-free mechanism would be an auction with the set of rules that do not depend neither on the characteristics of an object to be sold nor of the bidders (see Footnote 3). An English auction is a detail-free mechanism. An efficient multi-object mechanism of Dasgupta and Maskin (2000) is an example of a designed detail-free mechanism.

The second part of the Wilson's doctrine calls on the reduction of the set of common knowledge requirements imposed for the analysis.[21] Dominant strategies equilibria as well as ex post equilibria are particularly appealing concepts since for these equilibria little is required to be known by players about other players. In particular, players can be completely ignorant of how all players' types are distributed.

[21] Wilson (1987) also writes, "Game theory has a great advantage in explicitly analyzing the consequences of trading rules that presumably are really common knowledge; it is deficient to the extent it assumes other features to be common knowledge, such as one agent's probability assessment about another's preferences or information. I foresee the progress of game theory as depending on successive reductions in the base of common knowledge required to conduct useful analyses of practical problems. Only by repeated weakening of common knowledge assumptions will the theory approximate reality."

The problem is that these equilibria need not exist. And while they do exist in the ordinary English auction, adding an active seller invalidates stricter equilibrium constraints. Still, as we saw, in the dominant equilibrium buyers, but not the seller, have dominant strategies. Moreover, if we redefine dominance to account for the seller's ability to commit, then the buyers' strategies in both the optimal and the dominant equilibrium would be almost dominant (weakly better than any other strategy).

To account for the inherent asymmetries of the game considered and to be able to reduce common knowledge requirements, we propose two equilibrium concepts that impose different requirements on the seller and the buyers. For purposes of generality, let S_0 and S_i, $i = 1..N$ be sets of available strategies to the seller and the buyers correspondingly, $u_i(x_i, s_0, ..., s_n)$ for $i = 0..N$ are payoff functions.

Definition 5 *A* buyers' dominant strategy equilibrium *is a Bayesian-Nash equilibrium* $(s_0(x_0), s_1(x_1), ..., s_N(x_n))$ *such that for each buyer* i *of type* x_i *an additional constraint holds: for any* s_i' *and any* $\mathbf{s}'_{-i}$,

$$E_{x_0} u_i(x_i, s_0(x_0), s_i(x_i), \mathbf{s}'_{-i}) \geq E_{x_0} u_i(x_i, s_0(x_0), s_i', \mathbf{s}'_{-i}). \tag{6}$$

Definition 6 *A* buyers' ex post equilibrium *is a Bayesian-Nash equilibrium* $(s_0(x_0), s_1(x_1), ..., s_N(x_n))$ *such that for each buyer* i *of type* x_i *an additional constraint holds: for any* s_i',

$$E_{x_0} u_i(x_i, s_0(x_0), s_i(x_i), \mathbf{s}_{-i}(\mathbf{x}_{-i})) \geq E_{x_0} u_i(x_i, s_0(x_0), s_i', \mathbf{s}_{-i}(\mathbf{x}_{-i})).$$

These definitions are narrow since they are based on identities of the game considered in the paper. This is done to avoid excessive notation. The definitions can be easily generalized to any two disjoint subsets of players, with additional equilibrium conditions imposed only on one of the subsets. Note also, that commitment abilities of the seller (or of any player in general) are nowhere explicitly recognized. This can be done by changing the seller's constraint, requiring that she uses strategy $s_0^*(x_0)$ that maximizes expected payoff given all other players follow their best responses to $s_0^*(x_0)$ and to each others strategies (with additional constraints) without requiring $s_0^*(x_0)$ to be a best response. It turns out that in both presented equilibria, the seller's strategy is also a best response, so that at the moment the game is played it makes no difference whether the seller is able to commit.

Proposition 7 *Both, the dominant and the optimal equilibria are buyers' dominant strategies equilibria.*

Proof. It is straightforward. Note that, given that the seller's type (if not trivial) is immediately revealed, expectations over x_0 in (6) effectively disappear since buyers' strategies are conditional on x_0, and equilibrium conditions need to be checked for each given x_0. ■

What is the minimal amount of knowledge the players need to have for the presented equilibria to exist? Given that these equilibria are buyer's dominant strategy equilibria and in parallel with the properties of usual dominant strategies equilibria, the buyers need know nothing about distributions of the other buyers (and of the seller). The seller, naturally, needs to know everyone's distributions. In the dominant equilibrium buyers can be completely ignorant of their own distributions. Not so in the optimal equilibrium! Each buyer has to know her own distribution, and moreover it has to be common knowledge between her and the seller. This last feature is particularly interesting, since it suggests that the optimal equilibrium is not a robust equilibrium.[22]

It is interesting to note that in the detailed-free implementation of the optimal mechanism by Caillaud and Robert (2003) common knowledge requirements can be significantly reduced as well. There, the seller can be completely ignorant of the types' distributions. Each buyer's distribution needs to be a common knowledge between that and some other buyer.

4.4 Observable identities

So far we have assumed that the identities of the buyers and of the seller were observable to anyone. This is somewhat unrealistic assumption. A seller on eBay, for example, would make sure to conceal her identity. Moreover, a seller can potentially bid under several identities.

It turns out that our observable identities assumption can be significantly reduced while maintaining existence of both the dominant and the optimal equilibria. We can let the seller and the buyers bid under as many identities as she wishes, and assume that buyers cannot distinguish identity of any of the bidders. As an example of such auction, imagine each buyer bidding

[22] A *robust* equilibrium, as defined by Dasgupta and Maskin (2000), is a Bayesian Nash equilibrium in strategies that do not depend on types' distributions. In a way it is a different way to define an ex post equilibrium.

trough a representative, and that no buyer can recognize whom a particular bidder represents. What is important for the results, is that the seller can identify the identity of each buyer. The equilibria are then completed as follows. Each buyer would believe that the seller is active as long as at least two bidders are active and act accordingly.

What if the seller cannot identify the identity of each buyer? Or, in general, what if the seller does not know a distribution of values of a particular bidder, while knowing some aggregate information? An eBay is an example of such environment. Although the seller might know the identity of a particular bidder, it is more reasonable to think that she does not. At the same time, the seller might know that potential bidders can be of several types (experienced or not, e.g.) and have an idea of bid distributions of each type. Then, by observing bidding, the seller might receive additional information about participants' types and play accordingly. Graham, Marshall, and Richard (1990) and Graham, Marshall, and Richard (1996) use a setting like this. To provide a comprehensive analysis of the English auction with shill bidding, especially of its revenue properties is one of the avenues of the future research.

4.5 Interdependent valuations

A remarkable feature of an English auction is that, unlike other simple auction forms, it is efficient under quite general conditions in the interdependent values setting—when a buyer's private information may affect how other buyers value the object.[23] Moreover, this is done via an ex post equilibrium, independent of details of types' distributions.

We conjecture that results of this paper can be transferred to the interdependent values setting under appropriate conditions. In particular, it seems that if buyers valuations are separable in own signal, $V_i = x_i + \Psi_i(\mathbf{x}_{-i})$, and some other conditions on Ψ_i are imposed, then the optimal mechanism can be derived and implemented by the English auction via buyers' ex post equilibrium.

[23] See Milgrom and Weber (1982), Maskin (1992), Krishna (2003), Birulin and Izmalkov (2003), Izmalkov (2003). A mechanism is efficient if it has an equilibrium that allocates the object to the player who values it the most.

4.6 Auctions versus bargaining

Both presented equilibrium constructions require the seller to wait until only one buyer remains and then decide whether to shill bid—to make a take-it-or-leave-it offer or not. And, in general, whether the seller can commit or not, a strategy to wait until only one buyer remains seems to be "safe." The remaining stage of the game, the last active buyer versus the seller, looks like a bargaining game.[24] The outcome of it may as well depend on the participants relative bargaining powers.

If the seller can commit, which means she has a full bargaining power, she makes an optimal take-it-or-leave-it offer. When the seller has nor commitment power, as results of Vartiainen (2002) suggest, the maximal price the seller would get is the second highest value, that is the last remaining buyer gets all surplus from trade with the seller. Here is an equilibrium that supports this allocation (identity of the seller needs to be observable). Each buyer modifies the strategy from the dominant equilibrium as follows. Whenever a buyer and the seller remain the only active bidders, the buyer exits immediately. Then, if the seller decides to wait until only one buyer remains active, the seller runs a $1/2$ chance of retaining the object. A best response by the seller is to exit earlier (or not enter at all).

It would be interesting to further explore this bargaining-like feature of the English auction.

A Appendix

A.1 Shill bidding

A shill is a decoy or accomplice, especially one posing as an enthusiastic or successful customer to encourage other buyers, gamblers, etc., as defined in Oxford English Dictionary online. To shill is to boost for the auctioneer. Nowadays, especially with the development of online auctions, the usage of the term is much wider. The leading online marketplace, eBay, defines shill bidding as the deliberate placing of bids to artificially drive up the price of an

[24] Some real estate auctions have alike bargaining features. In addition to bids submitted, brokers on occasion pressure the highest bidder (and sometimes other as well) into raising her bid.

item. It is illegal and is punishable by suspension.[25] Bid padding, phantom bidding, bidding of the wall, lift-lining, trotting, running up the bid, setting hidden reserve prices—are other names to describe the same or very similar patterns of behavior both by sellers and auctioneers (see Cassady (1967) for an excellent account of various practices).

In general, mostly driven by online auctions, the term shill bidding is used to cover a variety of practices by the seller, ultimately to obtain more profit, but not necessarily by means of increasing the selling price directly by submitting bids through a shill, a friend or a different identity. A shill may also send bogus positive feedback to build or support a reputation as a reliable seller. The seller can also win the auction and then renege, with the purpose of exposing buyers' valuations and offering the buyer with the highest value another object exactly like the one just sold. Interestingly, eBay allows such contacts if the sale to the highest bidder cannot be completed or if the seller indeed is in possession of an identical object. Another way to expose potential winner value takes into account proxy bidding and works as follows. When a buyer, say *John*, submits a proxy bid, say of a $100 (presumably a true value), the current highest bid (and a provisional winning price) may show $50, meaning that the current second highest bid is $50 or just below. If somebody then submits a bid of $70, *John's* bid is automatically and in small increments increased to beat $70, so that the provisional winning price becomes $70 while the identity of the winner is not be changed. A seller's shill may submit an unreasonably high bid, say of $10000, which, at the moment of submission, raises the provisional price to $100 and reveals the true bid of *John*. The shill then may ask eBay to void her bid citing a bidding error. This would reduce the price back to $70. The same or another seller's shill may then submit a bid of $99, presumably extracting almost all surplus from *John*. In part, proponents of last minute bidding cite this practice as one of the reasons to bid at the very last moment, so that the seller would be unable to expose the true bid.

It is hard to provide any comprehensive account on how prevalent shill bidding is. The main reason is that it is considered to be illegal, and so sellers, when shill bidding, would try to conceal it. An interesting empirical study is done by Kauffman and Wood (2000) based on the data from 14,528 rare coin auctions on eBay during May 1999 and February 2000. They define

[25] From eBay's "Frequently asked questions about shill bidding," http://pages.ebay.com/help/basics/f-shilling.html.

questionable bidder behavior (QBB) as bidding on an item when the same or a lower bid could have been made on the exact same item in a concurrent auction ending before the bid-upon auction. *QBB*, presumably associated with sellers' shill bidding activities, was found in 987 bids in 713 auctions. While it is hard to say whether this is a lot or a little, the authors further study the bidders and sellers who are involved in *QBB* in greater detail. The authors identify four principles that a shill bidder would follow: "those who run up the bid: (a) are agents of the seller, and therefore not necessarily buyers and will tend to limit their bids to a single seller or perhaps a few seller Ids; (b) do not want to win the auction, but rather want the winner to pay more; (c) want to avoid bidding near the end of the auction where the chance of winning is greater; and, (d) bid in increments higher than average in an effort to quickly run up the bid." Then, the identify *questionable bidders (QB)*, those who exhibited *QBB*, and test the above principles. They found support for all four: *QBs* had more bids per seller than other bidders (1.45 vs. 1.25), indicating that *QBs* are concentrating on specific sellers (?); win only 26% while other bidders win 35% of the time; drop out on average 5.1 days before the auction ends compared with 1.8 day for the others; tend to bid 200% above the previous bid, if there is one, compared to 65% for the others.

References

AXELROD, R. (1984): *The Evolution of Cooperation.* Basic Books, New York.

BERGEMANN, D., AND S. MORRIS (2003): "Robust Mechanism Design," Cowles Foundation Discussion Paper No. 1421.

BIRULIN, O., AND S. IZMALKOV (2003): "On Efficiency of the English Auction," Penn State University and MIT, mimeo.

BULOW, J., AND J. ROBERTS (1989): "The Simple Economics of Optimal Auctions," *The Journal of Political Economy*, 97, 1060–1090.

CAILLAUD, B., AND J. ROBERT (2003): "Implementing the Optimal Auction," CIRANO, Scientific Series, working paper 2003s-31.

CASSADY, R. J. (1967): *Auctions and Auctioneering.* University of California Press, Berkeley and Los Angeles.

CHUNG, K.-S., AND J. C. ELY (2003): "Foundations of Dominant Strategy Mechanisms," Northwestern University, mimeo.

DASGUPTA, P., AND E. MASKIN (2000): "Efficient Auctions," *Quarterly Journal of Economics*, 115, 341–388.

GRAHAM, D. A., R. C. MARSHALL, AND J.-F. RICHARD (1990): "Phantom Bidding against Heterogeneous Bidders," *Economics Letters*, 32, 13–17.

——— (1996): "Lift-lining," in *Advances in Applied Microeconomics*, ed. by M. R. Baye, vol. 6. Jai Press Inc., Greenwich, Connecticut.

IZMALKOV, S. (2003): "English Auctions with Reentry," MIT, mimeo.

KATKAR, R., AND D. LUCKING-REILEY (2001): "Public Versus Secret Reserve Prices in eBay Auctions: Results from a Pokemon Field Experiment," NBER Working Paper No. 8183.

KAUFFMAN, R. J., AND C. A. WOOD (2000): "Running Up the Bid: Modeling Seller Opportunism in Internet Auctions," University of Minnesota, MISRC working paper 00-20.

KRISHNA, V. (2002): *Auction Theory*. Academic Press, San Diego, California.

——— (2003): "Asymmetric English Auctions," *Journal of Economic Theory*, 112, 261–288.

LOPOMO, G. (1998): "The English Auction is Optimal among Simple Sequential Auctions," *Journal of Economic Theory*, 82, 144–166.

MASKIN, E. (1992): "Auctions and Privatization," in *Privatization: Symposium in Honor of Herbert Giersch*, ed. by H. Siebert, pp. 115–136. Institut für Weltwirtschaft an der Univarsität Kiel.

MILGROM, P., AND R. WEBER (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122.

MYERSON, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

Riley, J. G., and W. F. Samuelson (1981): "Optimal Auctions," *American Economic Review*, 71, 381–392.

Skreta, V. (2003): "Optimal Auction Design under Non-Commitment," University of Minnesota, mimeo.

Vartiainen, H. (2002): "Auction Design without Commitment," Yrjö Jahnsson Foundation and the University of Helsinki, mimeo.

Wang, W., Z. Hidvegi, and A. B. Whinston (2001): "Shill Bidding in English Auctions," Emory University and the University of Texas at Austin, mimeo.

Wilson, R. (1987): "Game-theoretic Analyses of Trading Processes," in *Advances in Economic Theory, Fifth World Congress*, ed. by T. F. Bewley. Cambridge University Press, Cambridge, UK.