AUCTION DESIGN WITHOUT COMMITMENT

Hannu Vartiainen

HECER and University of Helsinki

Abstract

In anonymous platforms like the Internet, committing to honor the outcome of an auction is difficult since the seller can benefit by reauctioning the good. We argue that how information is processed within the auction mechanism is crucially important in such circumstances. In our model, the seller uses an intermediary to extract information from the buyers but is not tied to sell the good with terms that the mechanism proposes. Instead, she may reauction the good again via some other intermediary. There are no restrictions on how many times and through which mechanisms the good can be reauctioned. The buyers may also choose their outside option at any stage of the game. We argue that a sequentially rational seller can only implement a version of the English auction, in particular the popular version where bidders employ proxy bids. This is a consequence of the informational properties of the English action: it reveals just the right information for the seller to be able to commit to the mechanism. (JEL: C72, D44, D78)

1. Introduction

That the rules of an auction mechanism are not altered—and that the bidders trust this being the case—is a central prerequisite for the mechanism to work properly. For the seller, who is usually in the position of designing the mechanism, this requires a variety of commitments regarding communication, allocation, and payment processes. Commitment is especially important for the seller since, as demonstrated by McAdams and Schwarz (2007a, b), it is typically she who bears the costs associated with rulebreaking; forward-looking buyers anticipate any ex-post redesign of the mechanism, and adjust their play accordingly at the interim stage. This commitment problem can be seen as the reason for the market for intermediaries. Their business is to help the seller commit to the rules of the game.

The editor in charge of this paper was George Marios Angeletos.

Acknowledgments: I am indebted to three referees and George-Marios Angeletos, the editor, for comments and suggestions that significantly improved the paper. I also thank Philippe Jehiel, Klaus Kultti, Bart Lipman, Rafael Repullo, Hamid Sabourian, Hannu Salonen, Antti-Jussi Suominen, Juuso Välimäki, and Asher Wolinsky for important feedback, discussions, and encouragement in various stages of the project. Hospitality of the Center of Economics at the University of Munich, where part of this work was conducted, is gratefully acknowledged.

E-mail: hannu.vartiainen@helsinki.fi

But in anonymous platforms like the Internet, intermediaries or third parties may only have limited capacity to prevent parties from violating the rules of the game. For example, the seller may cast shill bids or the buyers may shade their bids, there may be ex-post bargaining over the good, etc. (see Boyd and Mayo 2000).¹ The most obvious way for the seller to change the rules is not to honor the mechanism once it has been played and the outcome has been determined. Instead, she may reauction the good via another mechanism.

The job of online auctions such as eBay, u-Bid, or Amazon is to determine the winner and the price of the auction on the basis of bids, that is, to serve as an *infomediary*—a device through which communication takes place.² However, they cannot enforce the trade.³ From the theoretical point of view this is a problem since the seller does have an incentive to reauction the good via another mechanism after the mechanism has been run, in order to extract further surplus from the winner. A curious fact is that this appears to happen very rarely.⁴ We argue that the reason why the online auctions can be credibly committed to is the way information is being processed within them.⁵

We study auction mechanisms that can be implemented when the seller can change the rules of the game by reauctioning the good once the auction mechanism has been executed but before the physical transaction has taken place. No restrictions are put on how many times the good can be reauctioned nor what the structure of a new auction mechanism may be. Also the buyers can leave the game whenever they want. We show that the way online auctions are designed can be seen as a rational response to the commitment problem—it allows parties to commit credibly to the original mechanism.

The most common online auction mechanism is the *ascending English auction* or its versions (by Lucking-Reiley, 2000a, they cover almost 90% of the online auctions). The popularity of the English auction is striking since it is not, in general, among the most profitable auction mechanisms (Myerson 1981).⁶ The variant that almost all the auction sites (including the aforementioned ones) encourage the sellers to use is the one where bidders post a confidential proxy or maximum bid and the winning bidder (bidder with the highest proxy bid) pays the price of the second-highest bid (plus a bidding increment). This implementation of the English auction reduces bidding costs as the bidders do not have to stay online during the time the auction runs. From

^{1.} For an extensive literature on shill bidding, see Chakraborty and Kosmopoulou (2004), McAdams and Schwarz (2007b), or Izmalkov (2005), and references therein.

^{2.} For surveys, see Lucking-Reiley (2000a) and Ockenfels, Reiley, and Abdolkarim (2007).

^{3.} For example, eBay is explicit in not taking measures against a seller who refuses to honor the transaction (see http://pages.ebay.com/help/buy/report-trading.html).

^{4.} According to eBay, less than 1% of auctions are filed for fraudulent behavior (see also Adams 2005). The vast majority of filed cases concern the failure to transfer the merchandise or the payment, or the misrepresentation of the quality of the good. In fact, refusal to honor the terms of trade (or equivalent) is not even represented in the Federal Trade Commission's classification of online auction frauds (see http://www.ftc.gov/bcp/edu/pubs/consumer/tech/tec07.shtm.

^{5.} I am indebted to a referee who suggested this interpretation.

^{6.} However, the English auction may be optimal in a restricted class of mechanisms, see for example Milgrom and Weber (1982) or Lopomo (1998, 2001).

the viewpoint of the bidders, the proxy ascending auction is equivalent to the second price (Vickrey) auction as they cast a single bid and the winner pays the second highest bid. However, from the viewpoint of the seller the two mechanisms differ as the proxy ascending auction does *not* reveal the value of the winning bid. The only publicly revealed pieces of information are (i) the identity of the winner, and (ii) the value of the second highest bid.⁷

We shall argue that for the seller to be able to commit to the mechanism, it is crucial that the auction mechanism does not reveal too much information. Our main result is that, in the absence of commitment, the *only* mechanism that can be credibly implemented is the English auction, or one of its versions (e.g., the proxy bidding mechanism).

An eloquent example of why too much information may render an auction mechanism unworkable can be found in Lucking-Reiley (2000b). He describes how the operation of the Vickrey auction by a stamp auctioneer becomes jeopardized as the auctioneer is not able to commit to the rules of the game. Once the value of the highest bid is known, the seller faces an irresistible temptation to raise the price by pretending that another bid was received just under the maximum amount. Of course, knowing this, the bidders are reluctant to reveal their true willingness to pay.⁸

Note that the English auction would *not* have been subject to the same concerns since it would only have revealed the second highest bid to the stamp auctioneer. To understand why the English auction can be committed to, a little more machinery is needed.

To model the idea that the parties do not have commitment power, we assume that there are no physical costs of changing the rules of the game, and that the rules can be changed any number of times. To operationalize this idea, we decompose a mechanism into two parts, an *information processing device* and an *implementation device*. The information processing device can be interpreted as a trusted infomediary like an online auction site. Its role is to transform reliably the buyers' messages to an output—a public signal. For example, in the first-price auction the profile of bids serves as the public signal. No technological constraints are imposed on the form of the information processing device. The task of the implementation device is to implement a physical outcome contingent on the signal or "recommendation" generated by the information-processing device. In the context of the first-price auction, the task of the implementation device is to sell the good to the highest bidder with the price of his bid.

The key assumption we make is that neither the bidders nor the seller can commit to the implementation device. That is, ex post, after the information processing device has done its job and the signal has been produced, the seller can turn to another (composite) mechanism rather than implement the outcome suggested by the implementation device. The bidders are always free to choose an outside option rather than accept the outcome of the implementation device.

^{7.} Many auction sites, including eBay, also make the other losing bids publicly verifiable.

^{8.} See Rothkopf, Teisberg, and Kahn (1990) for a similar argument.

Our aim is to characterize mechanisms that the seller does not want to redesign that is, mechanisms whose implementation device the seller wants to obey. To this end, we capture the seller's mechanism selection behavior by a rule σ that identifies, for each probability distribution p over the buyers' valuations, a mechanism $\sigma[p]$ that the seller chooses to implement under the belief p. The rule σ specifies the dynamics of the play by specifying how the changing beliefs affect the seller's behavior.

Incentive compatibility requires that the seller can commit to implement $\sigma[p]$ under belief p. Two conditions are imposed on σ that guarantee that a sequentially rational seller can indeed do this. The first is that the rule has to be *internally consistent*: selecting $\sigma[p]$ should not be in conflict with obeying σ later when information is generated within $\sigma[p]$ and p is updated accordingly. The second condition is that the rule must be *optimal*: under any belief p the seller should not be able to profit by implementing some other incentive feasible mechanism than $\sigma[p]$ in the class of mechanisms that she can commit to, given that she obeys σ in the future when p is affected.⁹ The latter property is dubbed as the *one-deviation property*. A stationarity condition, which requires that the rule is not conditioned on payoff irrelevant information, is also assumed.¹⁰

It should be emphasized that the methods developed in this work do not attempt to challenge or provide an alternative for the standard equilibrium techniques. Rather, the modeling here is meant to be consistent with them. The motivation for the reducedform modeling approach is expositional. Cleaning away the details of the extensive form makes the model simple and transparent, and allows us to focus on the aspects of strategic interaction that are central.

Our main result is that the payoff and information structure of any feasible mechanism—that is, a mechanism chosen by a mechanism selection rule that meets consistency, the one-deviation property, and stationarity—is a version of the English auction (indexed by a tie-breaking rule). Conversely, the rule that always chooses a version of the English auction is consistent, stationary, and meets the one-deviation property. Thus the English auction or its versions are essentially the *only* mechanisms that the seller can credibly implement without commitment.¹¹

Our argument is closely related to the famous *Coase conjecture* which says that in the one-buyer scenario the seller without commitment power is forced to sell the good with the price equal to the least possible valuation of the buyer (see Gul, Sonnenschei, and Wilson 1986; Fudenberg, Levine, and Tirole 1985; Ausubel, Cramton, and Deneckere 2002). The English auction is robust against commitment

^{9.} Incentive feasibility here requires that the mechanism is incentive compatible and also guarantees the bidders their outside option payoff under nontruthful messages. This property, which is called veto-incentive compatibility, is a necessary and sufficient condition for the bidders to play the game honestly if they can walk away at any stage of the game.

^{10.} Stationarity, which is important for the uniqueness result, can be defined in many ways. The key aspect of any definition is that the seller's behavior should not be dependent on unnecessary details. This could be motivated by a a desire to avoid computational burden: if the seller has been programmed to implement a mechanism under current belief, and an extra piece of information does not allow her to implement any more profitable mechanism, then her choice remains unchanged.

^{11. &}quot;Essentially" here means that the auction may also reveal some immaterial information.

problems for an analogous reason. The English auction reveals (i) the buyer with the highest valuation (the winner), and (ii) the valuations of all but the winner. Since the winner is known to have the highest valuation once the output has materialized, the seller cannot commit to sell the good to anyone but the winner with the price equal to the lower bound of his possible valuations—that is, the second-highest valuation. Hence, as the seller cannot commit to change the English auction, she can commit to implementing it. But this implies that the seller cannot commit to any action that could be improved ex post by running the English auction after the outcome has been determined. We show that this constraint is actually very severe: only the English auction itself satisfies it.¹²

An important limitation of our model is that, since the signals of the information processing device are public, private communication between the bidders and the seller is ruled out—that is, the case where the information processing device generates different signals to the seller and the bidders (Skreta 2010 is, to our knowledge, the only paper in the literature that allows private communication). There are three justifications for this assumption. Firstly, online auctions in the real world explicitly ban all private communication between the seller and individual buyers. Secondly, the assumption allows us to avoid the vexed problem of mechanism design under an informed principal. Thirdly, as private communication is likely to make commitment more difficult to achieve (more detailed information than necessary of the buyers' willingness to pay increases the profitability of the redesign), the seller would rather employ an intermediary that prevents private signals since the costs from such activity tend to be borne by the seller, as demonstrated by McAdams and Schwarz (2007a, b). They show, in particular, that when the seller employing the first-price auction cannot commit not to take further offers, trade may eventually take place after a long delay and with substantial delay cost. Importantly, when the delay cost is low, the resulting mechanism closely resembles the English auction. This is consistent with our model which predicts that only the English auction can be committed to.

Relatedly, it is important that the seller can commit to the information processing device or, equivalently, that the infomediary can commit not to leak information to the seller. As McAdams and Schwarz (2007a, b) argue, this commitment ability can be motivated by the idea that, since the intermediary does not own the object, he should have weak marginal incentives to get a higher price. Moreover, since intermediaries are typically long-run players in the market, reputational concerns are more likely to provide them with strong enough incentives not to break the commitment.

Commitment to the information processing device is not a problem when the mechanism can be interpreted as a cheap talk game which does not require private information processing. For example, the English auction can be implemented via cheap talk. See Krishna (2007), and the references therein, for conditions where this is possible.

^{12.} For studies on the no-gap case in the durable good monopoly scenario, see Ausubel and Deneckere 1989a, b.

This work adds to the literature on mechanism design without commitment by developing a conceptual framework in which commitment to mechanisms can be analyzed in a reduced form, without making detailed assumptions concerning the extensive form game. This simplifies the analysis greatly, and allows us to circumvent the need to restrict the way the mechanism can be designed or redesigned (apart from allowing private communication between the seller and the bidders).¹³

McAfee and Vincent (1997) study a more structured setup in which the seller can set a positive reserve price but cannot commit to not re-auctioning the good if the current bids do not exceed the reserve price. Assuming a fixed auction mechanism, they demonstrate that as the lag before potential reauction becomes short, the sequentially optimal (given reauctioning) reserve price produces the same expected revenue as an auction with a reserve price equal to the seller's valuation of the good.

Bester and Strausz (2001) and Skreta (2006) study commitment in the single-agent case where the designer can use a general communication device to extract information from the agent. However, the number of rounds is bounded (two in the case of Bester and Strausz 2001), which allows the problem to be solved backwards. The result found by Bester and Strausz (2001) is that in the one-buyer case the best mechanism is still direct (however, Bester and Strausz, 2000, show that with more than one agent this no longer holds). But, as opposed to the revelation principle, contracts are no longer fully revealing since, in equilibrium, the agent randomizes. Skreta (2006) shows that in a multistage bargaining game it is still optimal for the seller to post prices in each period rather than extract information via some complicated mechanism.

Skreta (2010) develops techniques for analyzing multiagent mechanism design problems without commitment. In her framework, the seller can reauction the good if it has not been previously sold. Unlike in McAfee and Vincent (1997), the seller may now employ any mechanism to do this. There are two key aspects which make Skreta (2010) and this paper quite different. First, Skreta (2010) focuses on cases in which the number of redesign rounds is bounded (one in most of the paper). Due to discounting, the basic tradeoff is between selling fast and reaching the final period, in which the seller can commit to the optimal mechanism. In our framework, where there is no cost of redesigning the mechanism, Skreta's (2010) problem would be vacuous as the seller would choose a null mechanism in all but the last period and then implement the optimal one. The second and more important difference is that Skreta (2010) allows private communication, which is a major complication. The main result of the paper is that despite the possibility of private communication, an optimal mechanism is simple and only uses public communication.

The problem of redesign is also akin to the literature on resale in auctions. Much of the focus in this literature has been on identifying the optimal auction with resale.¹⁴ However, this literature is fundamentally different from the current approach in that once a good is sold, the buyer who obtains the good—and becomes the new seller—is

^{13.} Calzolari and Pavan (2006b) study the alternative case: optimal information revelation along the contracting process when the principal's commitment is perfect.

^{14.} See Zheng 2002; Calzolari and Pavan 2006a; Garratt and Troger 2006; Hafalir and Krishna 2008.

privately informed about his willingness to sell the good. This, in general, jeopardizes efficiency in the after market (due to Myerson and Satterthwaite 1983). Thus the problem no longer has the recursive structure that drives the analysis of this paper; that the design problem and *re*design problem are conceptually similar and that they should be solved by using the same principles.

Further connections to the literature are discussed in the final section.

The paper is organized as follows. Section 2 specifies the setup and the game. Section 3 defines the solution. The results are stated in Section 4. Sections 5 and 6 conclude with discussion on the solution concept.

2. Setup

There is a seller of a single indivisible good and a set $N = \{1, ..., n\}$ of buyers. Seller's publicly known valuation of the good is 0. Buyer *i*'s privately known valuation θ_i is drawn from a discrete set $\Theta_i \subseteq \mathbb{R}_+$. Write $\Theta = \times_{i \in N} \Theta_i$ with a typical element $\theta = (\theta_i)_{i \in N}$, and $\Theta_{-i} = \times_{j \neq i} \Theta_i$ with a typical element $\theta_{-i} = (\theta_j)_{j \neq i}$.¹⁵ Denote by $\Delta(\Theta)$ the set of probability distributions *p* over Θ , and by p_i the *i*th marginal distribution of *p*.¹⁶

The set of allocations of the good is

$$A = \{(a_1, \ldots, a_n) \in \{0, 1\}^n : a_1 + \cdots + a_n \le 1\},\$$

where $a_i = 1$ if the good is allocated to *i* and $a_i = 0$ otherwise. Write $a = (a_1, \ldots, a_n)$. A money transfer from buyer *i* to the seller is denoted by $m_i \in \mathbb{R}_+$ and $m = (m_1, \ldots, m_n)$ is a profile of transfers. The set of all outcomes x = (a, m) is then $X = A \times \mathbb{R}_+^n$.

Now we define a *mechanism*. A mechanism does two things: processes information and implements an outcome. We separate these tasks. Denote by $\Delta(X)$ the set of probability distributions over X. A mechanism φ is a composite function

$$\varphi = g \circ h : \Theta \to \Delta(X),$$

consisting of an information processing device h and an implementation device g such that

$$h: \Theta \to \Delta(S)$$
 and $g: S \to X$,

where $\Delta(S)$ is the set of probability distributions over *S*, an open subset of an Euclidean space. That is, the information processing device *h* generates, after receiving the buyers' messages, a public signal $s \in S$. The signal *s* is the only information anyone—including the seller—obtains from *h*. The outcome function *g* then implements an

^{15.} That is, $p_i(\theta_i) = \sum_{\theta_i} p(\theta_i, \theta_{-i})$.

^{16.} Hence countable and without accumulation points. This assumption is for technical simplicity.

outcome $x \in X$ conditional on the realized signal *s*. That the implementation device is deterministic reflects the idea that the seller cannot make partial commitment, for example in the probabilistic sense, to implement an outcome before it is actually implemented.

Letting $H = \{h : \Theta \to \Delta(S)\}$ and $G = \{g : S \to X\}$ denote the sets of information processing devices and implementation devices, respectively, the set of all composite mechanisms is

$$\varphi = \{g \circ h : \Theta \to \Delta(X)\}$$

such that $g \in G$ and $h \in H$.

The support of distribution p is denoted by $\operatorname{supp}(p)$. Also write $h(\theta) = \{s : h(s : \theta) > 0\}$ and $h(\operatorname{supp}(p)) = \{s : h(s : \theta) > 0 \text{ and } \theta \in \operatorname{supp}(p)\}$. Given p, a signal $s \in h(\operatorname{supp}(p))$ of the information processing device h induces a posterior

$$p(\theta:s,h) = p(\theta)h(s:\theta) \bigg/ \sum_{\theta \in \Theta} p(\theta)h(s:\theta).$$

To economize on notation, write $p(\cdot : s, h) = p(s, h)$. Since the signals are public, $p(s, h) \in \Delta(\Theta)$ for all $s \in h(\operatorname{supp}(p))$. By the definition of the support, $\operatorname{supp}(p(s, h)) \subseteq \operatorname{supp}(p)$ for all h and for all s.

The mechanism $g \circ h$ is *constant under p* if h(supp(p)) is singleton. A constant mechanism implementing outcome x is denoted by

$$\mathbf{1}_x \in \varphi$$

A constant mechanism does not affect the beliefs and implements the same outcome with probability one. The two mechanisms $(g \circ h)$ and $(g' \circ h')$ are *outcome equivalent* under p if they induce the same outcome function: $(g \circ h)(\theta) = (g' \circ h')(\theta)$, for all $\theta \in \text{supp}(p)$. Finally, if the information provided by the mechanism $g \circ h$ is not finer than what is necessary to implement the outcome. That is, if g(s) = g(s') implies s = s' for all $s, s' \in h(\text{supp}(p))$, then we may write $p(s, h) = p(g(s), g \circ h)$.

Buyer θ_i 's and the seller's payoffs from the allocation x = (a, m) are, respectively,

,

$$u_i(a, m, \theta_i) = \theta_i a_i - m_i$$
$$v(a, m) = \sum_{i \in N} m_i.$$

Abusing the notation slightly, we may denote the payoffs of buyer θ_i and the seller from the allocation x = (a, m) by $u_i(x, \theta_i)$ and v(x), respectively.

3. Solution

The seller's problem is that she cannot commit to the implementation device g once the signal s has been produced by the information processing device h. Rather, she may

be tempted to design a new mechanism under her post-signal belief. In this section, we identify conditions that the mechanism needs to satisfy for the seller to credibly commit to it.

What makes the problem challenging is that we cannot employ backwards induction or a related structure since there is no final stage from which to start the recursion. That is, for any past history of redesigned mechanisms, there still exists a chance to redesign the current mechanism. That is why the solution has to be based on a "fixed point" argument.

In other words, whether the seller can commit to a mechanism depends on what mechanisms are available to her ex post, given the post-signal beliefs. But since ex post she can only select from mechanisms that she can commit to—the same question she faced at the ex-ante stage—the mechanisms that the seller can commit to need to be identified for all beliefs simultaneously. At the same time, consistency across ex-ante and ex-post beliefs and the seller's incentives must be honored.

We solve the mechanisms that can be committed to in two nested parts. First we specify conditions under which, by the revelation principle, the buyers could commit to a direct mechanism. Then we identify conditions under which the seller can commit to a direct mechanism given that the buyers can. This requires defining which mechanism the seller would implement under different (posterior) beliefs, if they were to materialize at the ex-post stage of a mechanism. Otherwise one cannot guarantee that the seller can commit to the mechanism in the first place.

Buyers' Incentives. We assume that the buyers can exit any point of the game. Thus any implementable mechanism $g \circ h$ must be *ex-post individually rational* (EXP-IR):¹⁷

$$u_i(g(s), \theta_i) \ge 0$$
,

for all $s \in h(\theta)$, for all $\theta \in \text{supp}(p)$, for all $i \in N$. Given p, buyer θ_i 's *interim* payoff from a mechanism $g \circ h$ is

$$\sum_{\theta_{-i}}\sum_{s}p(\theta)u_i(g(s),\theta_i)h(s:\theta).$$

By the *revelation principle* (Myerson 1979), an implementable mechanism must be *incentive compatible*. A mechanism $g \circ h$ is incentive compatible(IC) if

$$\sum_{\theta_{-i}}\sum_{s}p(\theta)u_i(g(s),\theta_i)[h(s:\theta)-h(s:\theta_{-i},\theta_i')] \ge 0$$

for all $\theta_i, \theta'_i \in \Theta_i$ and all $i \in N$. However, incentive compatibility and ex-post individual rationality are not independent conditions: the right of veto might be exercised at the off-equilibrium histories. The following simple extension of incentive

^{17.} *Interim* individual rationality requires that participation be weakly profitable before the output has been realized. Ex-post constraint has been analyzed by, for example, Forges (1993, 1998) and Gresik (1991, 1996).

compatibility resolves the problem by allowing i to veto the outcome even after his untruthful announcements.¹⁸

DEFINITION 1 (VETO-IC). Given p, a mechanism $g \circ h \in \varphi$ is *veto-incentive compatible* if

$$\sum_{\theta_{-i}}\sum_{s}p(\theta)[u_i(g(s),\theta_i)h(s:\theta) - \max\{u_i(g(s),\theta_i),0\}h(s:\theta_{-i},\theta_i')] \ge 0, \quad (1)$$

for all $\theta_i, \theta'_i \in \Theta_i$, for all $i \in N$.

Veto-incentive compatibility requires that truthful reporting forms a Bayes–Nash equilibrium even if vetoing is possible after an untruthful announcement. Any implementable mechanism must thus be veto-incentive compatible. For any p, denote the set of veto-incentive compatible mechanisms by

$$VIC[p] \subset \varphi$$
.

It is easy to see that any veto-incentive compatible mechanism is incentive compatible and ex-post individually rational (but not vice versa).¹⁹

Truthful announcements form a Bayes–Nash equilibrium in a veto-incentive compatible mechanism $\varphi = g \circ h$ *if* the seller can *commit* to following g after h has performed its information processing task—that is, has produced its signal s. Thus a mechanism maximizing the seller's payoff subject to veto-incentive compatibility could be interpreted as the seller's full commitment benchmark. Since veto-incentive compatibility concerns only the payoffs, any signal structure—even one that fully reveals the buyers' types—is consistent with veto-incentive compatibility. However, while signals do not affect anyone's payoff directly, they may do so indirectly, via the seller's behavior at the ex-post stage.

Seller's Incentives. The seller's expected payoff from the mechanism $\varphi = g \circ h$ is

$$V(\varphi, p) = \sum_{\theta} \sum_{s} p(\theta) v(g(s)) h(g(s) : \theta).$$

She wants to maximize her expected payoff subject to the constraint of not redesigning the mechanism after observing the signal *s* from the information processing device *h*. That is, of replacing the outcome g(s) with *another* mechanism in φ that generates her a higher expected payoff than g(s). Our task is to identify the conditions under which she will not do that.

Let the seller's (pure) mechanism design strategy be captured by a *choice rule* σ that specifies, for each prior belief p, the seller's choice of the mechanism under these beliefs. Since it is without loss of generality to focus on mechanisms that the buyers can commit to, the choice rule is assumed to satisfy

^{18.} Veto-incentive compatibility is due to Forges (1998), and is closely related to IC* of Matthews and Postlewaite (1989).

^{19.} Choose $\theta_i = \theta'_i$ in (1). We only need EXP-IR and IC in the remainder of the paper.

$$\sigma: \Delta(\Theta) \to \varphi \tag{2}$$

such that $\sigma[p] \in VIC[p]$, for all p. Then $\sigma[p]$ represents the mechanism that the seller implements under p. Rule σ in turn represents the dynamic mechanism selection strategy of the seller.

We now identify properties that the choice rule σ should satisfy. We argue that the sequential rationality of the seller, and the buyers' knowledge of this, requires that σ reflect internal consistency and optimization. For these conditions, we need to develop some concepts. We say that a mechanism $g \circ h \in \varphi$ is (weakly) *expost dominated* by a mechanism $\varphi \in \varphi$ if there is a signal $s \in h(\operatorname{supp}(p))$ such that

$$V(\varphi, p(s, h)) \ge v(g(s))$$
 and $\varphi \ne \mathbf{1}_{g(s)}$

under p(s, h). That is, the seller weakly prefers φ over the recommended outcome g(s), given the ex-post beliefs due to signal s. In such a case, the original mechanism $g \circ h$ may be subject to redesign. It is easy to see that for a typical p there is no veto-incentive compatible mechanism that is not ex-post dominated.²⁰ Thus the seller is typically (weakly) tempted to redesign the mechanism.

Now we define the set of mechanisms that the seller can commit to today given that σ is followed in the future. Under prior p, denote by $C^{\sigma}[p]$ the set of mechanisms that are *not* subject to redesign under the hypothesis that σ is followed ex post:

$$C^{\sigma}[p] = \left\{ g \circ h \in VIC[p] : \\ g \circ h \text{ is not ex-post dominated by } \sigma[p(s, h)], \text{ for any } s \in h(\operatorname{supp}(p)) \right\}.$$
(3)

Hence, by the revelation principle, and under the hypothesis that the seller can commit to the choice rule σ :

- A mechanism φ is truthfully playable if $\varphi \in C^{\sigma}[p]$, since then it will not be redesigned ex post.
- A mechanism φ is *not* truthfully playable if $\varphi \notin C^{\sigma}[p]$, since then it will be redesigned ex post.

Choice set $C^{\sigma}[p]$ is defined with respect to the assumed rule σ . We now formally specify conditions that sequential rationality imposes on the choice rule σ itself. The first condition requires consistency in the sense that employing σ *ex ante* should not contradict σ being employed ex post.

DEFINITION 2 (Consistency). Choice rule σ is *consistent* if $\sigma[p] \in C^{\sigma}[p]$, for all p.

The second condition implies optimality. Given σ and p, the seller should choose a mechanism that maximizes her payoff in the set $C^{\sigma}[p]$.

^{20.} If $0 \in \text{supp}(p_i)$ for all *i*, then a veto-incentive compatible mechanism is not ex-post dominated only if extracts all surplus from the buyers. But full surplus extraction á la Crémer and McLean (1984) is not possible under veto-incentive compatibility.

DEFINITION 3 (One-Deviation Property). Choice rule σ satisfies the *one-deviation* property if $V(\sigma[p], p) \ge V(\varphi, p)$, for all $\varphi \in C^{\sigma}[p]$, for all p.

Under the hypothesis that σ can be committed to in the future, the seller does not want to change σ under any current prior *p*. Without the one-deviation property, σ could not be convincingly committed to.

Note how consistency, the one-deviation property, and ex-post dominance play different roles in the solution. The one-deviation property together with consistency reflects optimization: $\sigma[p]$ maximizes the seller's payoff in $C^{\sigma}[p]$. Ex-post dominance in turn guarantees that this act of optimization is consistent with farsightedness. That is, since σ is obeyed in the future, not being ex-post dominated with respect to $\sigma[\cdot]$ guarantees that a mechanism can be committed to. Indeed the only role of ex-post dominance is to test whether the seller can commit to a particular mechanism under the hypothesis that σ is followed in the future. In particular, the one-deviation property is *not* implied by ex-post dominance.

Now we state two straightforward but important implications of consistency and the one-deviation property. First, the seller always implements the outcome of a mechanism that she can commit to.

LEMMA 1. Let σ be consistent and satisfy the one-deviation property. Then $g \circ h \in C^{\sigma}[p]$ implies that $\sigma[p(s, h)] = \mathbf{1}_{g(s)}$, for all $s \in h(supp(p))$.

Proof. Take any $s \in h(\operatorname{supp}(p))$. By consistency, $g \circ h$ is not ex-post dominated by $\sigma[p(s, h)]$ under p. By the definition of ex-post dominance, $\mathbf{1}_{g(s)}$ is not ex-post dominated by $\sigma[p(s, h)]$. Hence either $v(g(s)) > V(\sigma[p(s, h)], p(s, h))$ or v(g(s)) = $V(\sigma[p(s, h)], p(s, h))$ and $\sigma[p(s, h)] = \mathbf{1}_{g(s)}$. By the one-deviation property, $V(\sigma[p(s, h)], p(s, h)) \ge v(g(s))$. Hence it must be the case that $\sigma[p(s, h)] = \mathbf{1}_{g(s)}$.

In particular, the choice rule σ is *idempotent* in the following sense: if $\sigma[p] = g \circ h$, then $\sigma[p(s, h)] = \mathbf{1}_{g(s)}$, for all $s \in h(\operatorname{supp}(p))$. That is, running σ twice rather than once will not affect the outcome.

Second, if the seller can commit to implementing an outcome, then that outcome must maximize her payoff in the class of individually rational outcomes.

LEMMA 2. Let σ satisfy the one-deviation property. Then $\sigma[p] = \mathbf{1}_x$ implies that $v(x) \ge v(y)$, for all $\mathbf{1}_y \in VIC[p]$.

We now check that our solution is consistent with the standard bargaining theory.

The Coase Conjecture. The Coase conjecture, which pertains to our n = 1 case, argues that when the seller is unable to commit not to sell the good, the buyer is able to extract all the surplus. That is, the outcome of the one-sided bargaining game is to sell the good with price $\underline{\theta}(p)$, the minimal possible valuation θ in the support of p. The Coase conjecture has been extensively studied in the noncooperative bargaining

literature, and verified in the so-called "gap" case $\underline{\theta}(p) > 0$ by for example Fudenberg, Levine, and Tirole (1985) and Gul, Sonnenschein, and Wilson (1986).

The next remark shows that the result can be derived also in our setup, without going into the details of the bargaining process. Thus consistency and the one-deviation property do capture the key aspects of sequential rationality.

REMARK 1 (Gap case). Let n = 1. Let σ be a consistent choice function meeting the one-deviation restriction. Then $\sigma[p] = \mathbf{1}_{(1,\theta(p))}$, for all p such that $\underline{\theta}(p) > 0$.

That is, any $\sigma[p]$ sells the good to the buyer with the price equal to his minimal possible valuation. To see this, note that by Lemma 1, $\sigma[p(s, h)] = \mathbf{1}_{g(s)}$, for all $s \in h(\operatorname{supp}(p))$. By Lemma 2, g(s) maximizes v in the class of constant, individually rational mechanisms under p(s, h). Since $\underline{\theta}(p(s, h)) > 0$ we have $g(s) = (1, \underline{\theta}(p(s, h)))$. But by imitating $\theta = \underline{\theta}(p) > 0$, any $\theta' \in \operatorname{supp}(p)$ can guarantee to be able to buy the good at price $\underline{\theta}(p)$. Hence by incentive compatibility, $g(h(\theta)) = (1, \theta(p))$, for all $\theta \in \operatorname{supp}(p)$.

However, it is also well known that in the "no gap" case, $\underline{\theta}(p) = 0$, other more complex equilibria can be constructed (see e.g. Ausubel and Deneckere 1989a). To avoid them, the literature often focuses on simple "stationary" equilibria (see e.g. Ausubel, Cramton, and Deneckere 2002).

The problem with multiplicity of (complex) solutions also applies in our case when $\underline{\theta}(p) = 0$. It can be shown that for any $\lambda \in \Theta$ there is a choice rule σ^{λ} that is consistent and meets the one-deviation property, and sells to types $\theta \ge \lambda$ and never to types $\theta < \lambda$ of the buyer given the prior p. However, all constructed σ^{λ} are complex, and require the seller to condition σ^{λ} on seemingly superficial information. To remove these complexities, our final restriction imposes a degree of simplicity on choice rules. It demands that the implemented outcome is not conditioned on information that does not provide more profitable transaction opportunities.

DEFINITION 4 (Stationarity). A choice rule σ is stationary if $\sigma[p] = \mathbf{1}_x$, $\sigma[p'] = \mathbf{1}_{x'}$, $v(x) \ge v(x')$, and $supp(p') \subseteq supp(p)$ imply x = x'.

That is, signals that do not allow the seller to implement a more profitable choice do not affect the seller's choice.²¹ For example, the choice rule σ^{λ} above fails stationarity (see Appendix B for the precise exposition). The next section characterizes the inducible stationary choice rules in the general $n \ge 1$ case.

4. Results

4.1. Overview

Before formally establishing the results, we give a heuristic account of them. Our main result is that the English auction (or its version) is the only mechanism that

^{21.} For an analogous restriction, see Ausubel and Deneckere 2002.

the seller can credibly commit to. The core of our argument is simple. The outcome of the English auction reveals exactly the information that is needed for the seller to credibly implement this outcome. The reason for this is that the outcome of the English auction separates the possible valuations of the winner from those of the other buyers, and hence the seller cannot commit to not selling to the winner with the proposed price—his least possible valuation given the ex-post information. We show that only the English auction has this property.

More specifically, we first demonstrate that the choice rule where the seller always uses the English auction—or the version which uses all the information contained in the nonwinners' valuations to compute the lower bound of the winner's valuation—is consistent. This follows from the fact that running the English auction twice rather than once does not affect the outcome. To prove that the English auction choice rule also satisfies the one-deviation property, we show that if a mechanism cannot be profitably changed by employing the English auction ex post, then the mechanism must, in fact, be outcome equivalent with the English auction itself. This "fixed point" result is obtained via induction on the value spaces of the buyers, and implies that the English auction choice rule cannot be changed profitably with a one-time deviation.

To prove that the English auction choice rule is the *only* feasible one, we first argue that any seller's choice rule that is stationary and consistent cannot reject (all the versions of) the English auction. This is true since any realization of an English auction reveals the winner and his least possible valuation given the other buyers' valuations. Thus any mechanism that ex-post dominates the English auction must threaten the winner to sell the good to the buyer with the second highest valuation, to force the winner to pay a higher price. However, this threat is not credible since when the winner declines the offer, the seller cannot commit to not sell—as suggested by the Coase conjecture—to the winner with his least possible valuation. Finally, since some English auction is always feasible, it follows that the mechanism chosen by the seller must be robust against being upset by the English auction ex post. This strong restriction implies, by the fixed point argument made in the previous paragraph, that any choice rule must be outcome equivalent with the English auction choice rule.

Two modeling aspects complicate the analysis. First, we do not put any restrictions, for example in terms of independence, on the distributions of the buyers' valuations. Doing otherwise would require us to impose unnatural constraints on the employed mechanisms since the posterior distributions should also satisfy the restrictions. To cope with general distributions, the English auction has to be modified to take into account all the information hidden in the nonwinners' valuations which makes the formal definition of the mechanism somewhat more involved than usually. Second, since we cannot assume away the positive probability of an event where at least two buyers have same highest valuation or the highest valuation equals the seller's valuation, the English auction has to be appended with a tie-breaking rule. By consistency, the same tie-breaking rule has to be applied uniformly to the whole English auction choice rule. Different tie-breaking rules then give rise to different English auction choice rules.

4.2. The Analysis

The English Auction φ^E . A tie-breaking rule $w : \Theta \to N \cup \{0\}$ always selects one of the players with the highest valuation

$$w(\theta) \in \arg \max_{j \in N \cup \{0\}} \theta_j, \tag{4}$$

for all $\theta \in \Theta$. Note that $w(\theta) = 0$ means that the seller keeps the good when $\theta_i = 0$ for all $i \in N$. The buyer $w(\theta)$ is called "the winner" under the type profile θ , and the set $w^{-1}(i) \subseteq \Theta$ consists of the type profiles under which w chooses i as the winner. Our mechanism is indexed by the tie-breaking rule.

Given a prior distribution p and a tie-breaking rule w, we now construct a deterministic mechanism $\varphi^E(\cdot) = g^E(h^E(\cdot))$, the *English auction*.²² Define a set of public signals S^E by

$$S^{E} = \{ (w(\theta), \theta_{-w(\theta)}) : \theta \in \Theta \},\$$

and a deterministic information processing device $h^E: \Theta \to S^E$ such that

$$h^{E}(\theta) = (w(\theta), \theta_{-w(\theta)}),$$

for all $\theta \in \Theta$. In words, under type profile θ the information processing device of the English auction produces the signal $h^{E}(\theta) = (w(\theta), \theta_{-w(\theta)})$, which makes public the identity of the winner $w(\theta)$ and all the nonwinners' types $\theta_{-w(\theta)}$.

The implementation device g^E is defined on S^E . The English auction allocates the good to the winner with the price equal to the least possible valuation of the winner above the second highest valuation. Formally, the winner's money transfer rule is as follows: for any $i \in N$,

$$m^{E}(i, \theta_{-i}, p) = \min\{\theta'_{i} : (\theta'_{i}, \theta_{-i}) \in w^{-1}(i) \cap \operatorname{supp}(p)\},$$
(5)

if $\theta \in w^{-1}(i)$. That is, $m^E(i, \theta_{-i}, p)$ is the least possible valuation of *i* given the information that (i) *i* is the winner, (ii) the other buyers' types are θ_{-i} . The implementation device $g^E : S^E \to X$ is now defined, for each buyer $j \in N$, for all $(i, \theta_{-i}) \in S^E$, by

$$g_{j}^{E}(i,\theta_{-i}) = \begin{cases} (1,m^{E}(i,\theta_{-i},p)), & \text{if } j = i, \\ (0,0), & \text{if } j \neq i. \end{cases}$$
(6)

That is, $\varphi^{E}(\theta) = g^{E}(h^{E}(\theta))$ allocates the good to the winner $i = w(\theta)$ who pays a price equal to his least possible valuation (i) given the other buyers' types, (ii) the fact that *i* is the winner, and (iii) p^{23} The corresponding payoffs are, for all $i \in N$,

$$u_i(\varphi^E(\theta), \theta_i) = \begin{cases} \theta_i - m^E(i, \theta_{-i}, p), & \text{if } \theta \in w^{-1}(i) \cap \operatorname{supp}(p), \\ 0, & \text{if } \theta \notin w^{-1}(i) \cap \operatorname{supp}(p). \end{cases}$$

^{22.} We relax w from the description of the English auction φ^E for notational simplicity.

^{23.} Note that $\varphi^{E}(p)$ reveals only the winner's identity and the other players' valuations. Hence it cannot be interpreted as the Vickrey (second-price) auction which asks *all* buyers to reveal their valuations.

By construction, φ^E is *efficient* and the price paid by the winner is less than or equal to his valuation, and at least as high as the other buyers' valuations.²⁴ Moreover, since the winner *i* becomes publicly known along with the signal $s = (i, \theta_{-i})$, the posterior belief $p((i, \theta_{-i}), h^E)$ satisfies

$$\operatorname{supp}(p((i, \theta_{-i}), h^E)) \subseteq w^{-1}(i).$$
(7)

Since the implementation device is a bijection, we may denote the posterior $p(s, h^E)$ simply by $p(x, \varphi^E)$.

Mechanism φ^E has the familiar pivotal structure: a buyer's payment—and hence his payoff—is independent of his announcement as long as he wins (or loses). The impact of lying on his payoff cannot be positive since it either induces the buyer to win when he would like to lose, or to lose when he would like to win, given the other buyers' valuations. Since truthtelling is a dominant strategy,

$$\varphi^E \in VIC[p].$$

Feasible Choice Rules. To highlight the fact that in the above definition the English auction is conditioned on the prior p, let us denote it by $\varphi^E[p]$. Keeping fixed the tiebreaking rule for which the English auction is defined, the function $\varphi^E : \Delta(\Theta) \rightarrow \varphi$ can now be taken as *the English auction choice rule*. Different tie-breaking rules then yield different English auction choice rules. Construct a correspondence $C^{\varphi^E} \subset VIC$ such that

 $C^{\varphi^{E}}[p] = \{\varphi \in VIC[p] : \varphi \text{ is not ex-post dominated by } \varphi^{E}[p'], \text{ for any } p' \in \Delta(\Theta)\}.$

Our aim is to prove that the English auction choice rule is essentially the only rule that can be committed to (up to a tie-breaking rule). The argument is made in two parts, the first demonstrating sufficiency and the second necessity. The result is stated in Theorem 1.

We first prove the sufficiency—that the English auction choice rule is consistent and satisfies the one-deviation property. This is done via two lemmas. First, note that the English auction choice rule, being defined with respect to a fixed tie-breaking rule, is idempotent: $\varphi^{E}[p(x, \varphi^{E}[p])] = \mathbf{1}_{x}$, for all $x \in \varphi^{E}[p](\operatorname{supp}(p))$, for all p. That is, after running the English auction, a new English auction does not change the outcome. This implies that the English auction $\varphi^{E}[p]$ is *not* ex-post dominated by $\varphi^{E}[p(x, \varphi^{E}[p])]$. Hence it follows that $\varphi^{E}[p] \in C^{\varphi^{E}}[p]$ for all p. More compactly: the English auction choice rule is consistent.

LEMMA 3. $\varphi^{E}[\cdot]$ is consistent.

We now argue that the English auction choice rule satisfies the one-deviation property. To do this, we establish a very strong result: *any* mechanism in $C^{\varphi^E}[p]$ is

^{24.} When the valuations are correlated, there may be a gap between this and the second-highest valuation.

outcome equivalent—and hence payoff equivalent—with the English auction $\varphi^{E}[p]$. This guarantees that there is no profitable single-stage deviation to the seller's mechanism design strategy. The result is also used in the proof of the uniqueness of the feasible seller's strategy.

LEMMA 4. $\varphi \in C^{\varphi^E}[p]$ only if φ is outcome equivalent to $\varphi^E[p]$, for all p.

Proof. Relegated to Appendix A.

That is, given p, the only veto-incentive compatible mechanisms that are not expost dominated by any $\varphi^{E}[p']$ are the mechanism $\varphi^{E}[p]$ itself and its versions that may additionally reveal some nonrelevant information concerning the winner's type. Thus φ^{E} has a "fixed point" property. The proof relies on induction on the buyers' type sets.

We now demonstrate that there is no other rule than the English auction choice rule that meets the desiderata. The argument is made in two stages. First we demonstrate that any stationary rule σ that is consistent and meets the one-deviation property induces C^{σ} that contains φ^{E} , specified for some tie-breaking rule w. The second part uses this finding to show that C^{σ} is contained in $C^{\varphi^{E}}$. Together with Lemma 4, this proves that φ^{E} is outcome equivalent to *any* feasible mechanism.

Heuristically, if φ ex-post dominates $\varphi^{E}[p]$, then φ must change $\varphi^{E}[p]$'s allocation. Since any outcome x of $\varphi^{E}[p]$ reveals the winner and his least possible valuation given the other buyers' valuations, φ must threaten the winner to sell the good to the buyer with the second-highest valuation to force the winner to pay a higher price. However, this threat is not credible since when the winner declines the offer, the seller sells—by stationarity and Lemma 2—to the winner with his least possible valuation.

LEMMA 5. Let a choice rule σ be stationary, consistent, and satisfy the one-deviation property. Then there is a tie-breaking rule wsuch that $\varphi^{E}[p] \in C^{\sigma}[p]$, for all p.

Proof. Construct *w* as follows: For any $\theta \in \Theta$, denote by $\mathbf{1}_{\theta}$ the degenerate prior such that supp $(\mathbf{1}_{\theta}) = \{\theta\}$. Then there is an outcome x_{θ} such that $\sigma[\mathbf{1}_{\theta}] = \mathbf{1}_{x_{\theta}}$. Let $w(\theta) = i$ if x_{θ} allocates the good to *i*. By Lemma 2, such $w(\theta)$ satisfies (4). Use this *w* to construct φ^{E} .

Suppose, in contrast to the claim, that there is p such that $\varphi^{E}[p] \notin C^{\sigma}[p]$. Then $\varphi^{E}[p]$ is ex-post dominated by $\sigma[p(x, \varphi^{E}[p])]$ for some $x \in \varphi^{E}[p](\operatorname{supp}(p))$. Denote $\sigma[p(x, \varphi^{E}[p])] = g \circ h$.

Let *x* allocate the good to player *i*. By stationarity, because $g \circ h$ ex-post dominates $\mathbf{1}_x$, $g \circ h$ cannot be a constant mechanism. By IC there are $\theta' \in \text{supp}(p(x, \varphi^E[p]))$, $s \in h(\theta')$, and $j \neq i$ such that $\theta'_i = \theta'_j$, and such that g(s) allocates the good to *j*. By Lemma 1, $\mathbf{1}_{g(s)} = \sigma[p(x, \varphi^E[p])(s, h)]$. Since

$$\mathbf{1}_{x} \in VIC[p(x, \varphi^{E}[p])] \subseteq VIC[p(x, \varphi^{E}[p])(s, h))],$$

it follows by Lemma 2 that $v(g(s)) \ge v(x)$.

By equation (6), $\sup(p(x, \varphi^E[p])) \subseteq Y_i$ and, by the definition of support,

$$\operatorname{supp}(p(x,\varphi^{E}[p])(s,h)) \subseteq \operatorname{supp}(p(x,\varphi^{E}[p])).$$

Thus, by the construction of w, EXP-IR, and Lemma 1, $\sigma[\mathbf{1}_{\theta'}] = \mathbf{1}_x$. However, since $g(s) \in h(\theta')$, also $\operatorname{supp}(\mathbf{1}_{\theta'}) \subseteq \operatorname{supp}(p(x, \varphi^E[p])(s, h))$. This implies, by stationarity, that g(s) = x, violating the assertion that x allocates the good to i and g(s) to $j \neq i$.

We now use the previous lemma to prove that if a stationary σ is consistent and meets the one-deviation property relative to *C*, then there is a tie-breaking rule *w* such that no element of C[p] is ex-post dominated by $\varphi^{E}[q]$ for any *q*.

LEMMA 6. Let a choice rule σ be stationary, consistent, and satisfy the one-deviation property. Then $C^{\sigma}[p] \subseteq C^{\varphi^{E}}[p]$, for all p, for some tie-breaking rule w.

Proof. Let $g \circ h \in C^{\sigma}[p]$ and $s \in h(\operatorname{supp}(p))$. Denote x = g(s) and q = p(s, h). By Lemma 1, $\sigma[q] = \mathbf{1}_x$. Identify w as in Lemma 5. It suffices for us to show that $\mathbf{1}_x$ is not ex-post dominated by $\varphi^E[q]$. Suppose, on the contrary, that it is. By Lemma 5, $\varphi^E[q] \in C^{\sigma}[q]$. By the definition of one-deviation property, $V(\sigma[q], q) \ge V(\varphi^E[q], q)$. Thus $v(x) \ge V(\varphi^E[q], q)$.

Take any $y \in \varphi^{E}[q](\operatorname{supp}(q))$. Then $\operatorname{supp}(q(y, \varphi^{E}[q])) \subseteq \operatorname{supp}(q)$ and, hence,

$$x \in \{x' : \mathbf{1}_{x'} \in VIC[q]\} \subseteq \{x' : \mathbf{1}_{x'} \in VIC[q(y, \varphi^E[p])]\}.$$

Since, by Lemma 1, $\sigma[q(y, \varphi^E[q])] = \mathbf{1}_y$ and $\mathbf{1}_x \in VIC[q(y, \varphi^E[q])]$, it follows by Lemma 2 that $v(x) \leq v(y)$. Since y was arbitrary, and $v(x) \geq v(\varphi^E[q], q)$, the inequality must hold as equality. But then, since $\operatorname{supp}(q(y, \varphi^E[q])) \subseteq \operatorname{supp}(p)$, stationarity implies that x = y. Thus $\varphi^E[q] = \mathbf{1}_x$, which contradicts the hypothesis that $\varphi^E[q]$ ex-post dominates $\mathbf{1}_x$.

By Lemma 6, a stationary and consistent σ that meets the one-deviation property is not ex-post dominated by an English auction. Hence $\sigma[p]$ cannot allocate the good to anyone but the buyer with the highest valuation—that is, $\sigma[p]$ must be efficient. An implication of the lemma is that the commitment inability of the seller leads to an efficient allocation, as suggested by the Coase theorem.

Since $\sigma[p]$ is efficient, and the lowest type of a buyer earns zero payoff, the *revenue* equivalence theorem implies that $\varphi^{E}[p]$ is the (generically) unique implementable mechanism if the buyers' valuations are independent. However, we can say more: by Lemma 4, if the seller is unable to commit, the uniqueness of the implementable mechanisms is a general phenomenon and holds for any prior distribution.

As an illustrative example, let $N = \{1, 2\}$ and $\Theta = \operatorname{supp}(p) = \{5, 10\}^2$. Let, say, w(10, 10) = w(10, 5) = w(5, 5) = 1 and w(5, 10) = 2. Take $\varphi \in C^{\sigma}[p]$. Since $\mathbf{1}_x$ is not ex-post dominated by $\varphi^E[p(x, \varphi)]$ for any $x \in \varphi(\operatorname{supp}(p)), \varphi(\theta)$ allocates the good to buyer 1 under all $\theta \in \{(5, 5), (10, 5), (10, 10)\}$. Transfers from 1 under $\theta = (5, 5)$ and $\theta = (10, 10)$ are 5 and 10, respectively. By incentive compatibility, transfer from 1 under $\varphi(10, 5)$ is 5. Since only $\varphi(5, 10)$ allocates the good to 2, type $\theta_2 = 10$ is then revealed. Hence, his transfer must be 10, which means that $\varphi = \varphi^E[p]$ under p. Now we are ready to state our main result. Since $\varphi^{E}[\cdot]$ is well defined (by equations (4), (5), and (6)), the first part of the theorem also proves the existence of a solution.

THEOREM 1. 1. Choice rule φ^E is stationary, consistent and satisfies the one-deviation property, for any tie-breaking rule w.

2. If a choice rule σ is stationary, consistent, and satisfies the one-deviation property, then there is a tie-breaking rule w such that $\sigma[p]$ is outcome equivalent to $\varphi^{E}[p]$, for all p.

Proof. 1. Stationary of φ^E follows from the existence of a tie-breaking rule w for which φ^E is constructed. Lemma 3 establishes consistency. The one-deviation property follows from Lemma 4: any $\varphi \in C^{\varphi^E}[p]$ agrees with $\varphi^E[p]$ on X and, hence, induces the same payoff as $\varphi^E[p]$.

2. Let stationary rule σ be consistent and satisfy the one-deviation property. By Lemma 5, there is w and C^{σ} such that $\varphi^{E} \in C^{\sigma}$. By Lemma 6, $C^{\sigma} \subseteq C^{\varphi^{E}}$. By construction, $\sigma \in C^{\sigma}$. Thus, by Lemma 4, $\sigma[p]$ is outcome equivalent to $\varphi^{E}[p]$, for all p.

That is, the seller can commit to the English auction provided that she does so consistently, under all distributions of the buyers' valuations. Moreover, the payoff structure of every feasible auction coincides with that of the English auction (defined for some tie-breaking rule w). The only difference of a committable mechanism and the English auction may concern additional, payoff irrelevant information on the player's valuations.²⁵

We should stress that the first result requires that the tie-breaking rule is well defined, and the second that the stationarity condition is satisfied by the choice rule. While the former assumption is innocent, and a consequence of appropriate modeling, the second one is more substantial. The key justification for stationarity is simplicity. A construction along the lines of Ausubel and Deneckere (1989) would presumably allow the seller to commit to a choice rule that is more complex than the English auction choice rule.

The Generalized Coase Conjecture. With the stationarity assumption the Coase conjecture can now also be verified in the "no gap" case. By the second part of Theorem 1, if σ is stationary, consistent, and meets the one-deviation property, then $\sigma[p] = \varphi^E[p]$ which allocates the good to any buyer $\theta > 0$ with price min $w^{-1}(1) \cap \operatorname{supp}(p)$.²⁶

To conclude, our argument can be seen as a generalization of the Coase conjecture which, in the current set up can be formulated as follows: $\varphi^{E}[p]$ is the unique feasible

^{25.} It is interesting that while full surplus extraction is feasible under full commitment under almost all p (Crémer and McLean 1988; McAfee and Reny 1991), only the English auction is feasible without commitment.

^{26.} If $\operatorname{supp}(p) > 0$, then this number is equal to $\operatorname{supp}(p)$. If $\operatorname{supp}(p) = 0$, then it is 0 if the tie-breaking rule allocates the good to the buyer at $\theta = 0$, and $\min\{\theta > 0 : \theta \in \operatorname{supp}(p)\}$ if it allocates the good to the seller at $\theta = 0$.

mechanism in the n = 1 case when the seller cannot commit to not sell to the buyer who values the good more than she does. Our more general version of the claim says that in the $\varphi^{E}[p]$ is the unique feasible mechanism in the $n \ge 1$ case when the seller cannot commit to not selling to the buyer who values the good more than the buyer with the second-highest valuation. More succinctly, *without external commitment devices*, *the seller can only commit to the English auction*.

5. Discussion of the Solution Concept

It is important to understand in what sense our key desiderata—the one-deviation property and consistency—reflect sequential rationality. Suppose that the seller does make a single deviation to her mechanism design strategy σ . Then under p she chooses another incentive feasible mechanism, say φ , rather than $\sigma[p]$, and returns to follow σ once φ has been processed and beliefs have been affected accordingly. We argue that the deviation cannot possibly be profitable for the seller. On the one hand, if following σ is consistent with the implementation of φ , then, by the one-deviation property and consistency of σ , φ could not have been a profitable deviation to $\sigma[p]$ in the first place. On the other hand, if following σ is *not* consistent with the implementation of φ , then φ will be redesigned expost according to σ and, since the deviation was of one-shot nature, the new mechanism is implemented as planned. But by the standard revelation argument, the combination of φ and its redesign can now be interpreted as a single incentive feasible mechanism that cannot be, by the one-deviation property and consistency of σ , a profitable deviation to $\sigma[p]$. This argument obviously generalizes to all finitely long deviations to the mechanism design strategy. Hence the one-deviation property and consistency of σ reflect sequential rationality of the players in the sense that no finite sequence of deviations to σ can be profitable for the seller.

If also infinitely long deviations are permitted, then an assumption has to be made concerning the payoffs from infinite streams of redesigns. There are natural assumptions that guarantee that the results of this paper still hold. Construct an explicit extensive form game in which the seller, at each round, designs an incentive feasible mechanism and then chooses whether to implement the outcome of the mechanism played by the bidders, or whether to design another incentive feasible mechanism to extract further information from the bidders. Also each infinite terminal history is now associated a payoff. We conjecture that if the payoffs are zero when the outcome is never implemented and the future is undiscounted, then the one-deviation property and consistency characterize the perfect Bayesian equilibria of the game as in this paper. If, however, the future is discounted or there are other redesign costs, then the equilibria of the game need not correspond to our equilibria (see Skreta 2010). Whether the equilibria do coincide in the limit, when the redesign costs tend to zero, remains an open question.

One may wonder whether the ex-post domination criterion in the definition of the one-deviation property, which is defined by a weak payoff dominance, is needlessly strong. A natural weaker candidate would be to demand strict payoff dominance.

Strict domination is, however, in conflict with our basic assumption that the seller's mechanism selection rule is dependent only on the prior p. To see this, consider the n = 2 case and $\text{supp}(p) = \{0\} \times \{0, 1\}$. With the tie-breaking rule w that allocates the good under $\theta = (0, 0)$ to buyer 2, mechanism $\varphi^E[p]$ would always sell to buyer 2 with price 0. With strict domination criterion, a procedure that sells to 1 under $\theta = (0, 0)$ with price 0 and to 2 under $\theta = (0, 1)$ with price 1 would be not be strictly ex-post dominated. And selling to 1 under $\theta = (0, 0)$ with price 0 would be in conflict with $\varphi^E[q]$ where prior q is degenerate on $\theta = (0, 0)$. Combining strict dominance with sequential rationality would therefore require history dependent choices, and the mechanism selection rule σ would no longer be definable as a function of p.

6. Conclusion

Mechanism design requires commitment since at the ex-post stage, when the mechanism has produced the information needed for implementing the output, the seller may want to change the rules of the game and design a new mechanism. Mechanisms are particularly vulnerable to commitment problems in anonymous platforms like the Internet.

We have studied auction mechanisms that the seller can credibly implement without external commitment devices. To this end, we have developed a new reducedform modeling approach. We recognize that a mechanism does two things: processes information and implements an outcome. We separate these tasks, and relate the latter as the source of the seller's commitment problem. If the seller cannot commit to implementing the outcome once the mechanism has been run, she has an incentive to extract further surplus from the buyers by initiating a new mechanism under the ex-post beliefs. Of course, the new mechanism is subject to the same commitment problem, and so on. The objective of the paper is to identify the conditions under which the seller will be able to commit to the mechanism design strategy has to be defined for all distributions. The two conditions we impose on the seller's mechanism design strategy, which reflect her sequential rationality, are internal consistency and optimization. We show that the only mechanism that satisfies these restrictions as well as a stationarity condition is the traditional English auction (or its version).

At the heart of the analysis is the argument that a sequentially rational seller can always commit to the English auction. This idea can be seen as a generalization of the Coase conjecture (e.g. Fudenberg, Levine, and Tirole 1985; Gul, Sonnenschein, and Wilson 1986). In the one-buyer case, the claim states that the seller can only commit to sell the good to the buyer with a price equal to his least possible valuation. In the multiple-buyers case, the seller can always commit to the English auction since, when the mechanism reveals the buyer with the highest valuation and his least possible valuation (= the second highest valuation), the seller cannot commit to raising the price above this lower bound. Importantly, because the seller *can* credibly commit to the English auction, she *cannot* commit to mechanisms that are ex-post dominated by

the English auction. Our main result is that this constraint is very severe: only the English auction itself satisfies it.²⁷

Our model provides support for the English auction in the *general class* of auction mechanisms. Many studies have demonstrated the usefulness of the English auction in a restricted class of mechanisms. In a classic treatise, Milgrom and Weber (1982) show that the English auction is optimal among the four standard auction forms when the valuations are affiliated, a natural assumption in many auction scenarios.²⁸ In the same model, Lopomo (1998, 2001) demonstrates that the English auction features *robustness* in a sense that it is optimal among simple sequential auctions and in a class of *posteriorly implementable* auctions (a concept due to Green and Laffont 1987).

Posterior implementability requires that the buyers' behavior is regret-free in the sense that they would not want to change their behavior even if they knew the outcome of the mechanism. This property is at the heart of the robustness of the English auction and the Vickrey auction.²⁹ It also partly drives our results. Due to posterior implementability, running the English or the Vickrey auction twice in a row rather than just once will not affect the outcome. But this is only a necessary condition of a feasible auction. On the sufficient side, one needs to guarantee that the seller cannot commit to any *more* profitable auction at the ex-post stage, given the information that is generated by the mechanism. For example, the Vickrey auction, which is posteriorly implementable, is not implementable when the seller cannot commit since it reveals too much information.

Finally, our model provides some insights into the literature on optimal auctions under efficiency (e.g. Ausubel and Cramton 1999; Krishna and Perry, 1998). The efficiency restriction is usually motivated vaguely by appealing to "Coasian dynamics", which leads to the efficient allocation of resources through the seller's inability to commit, or resale markets.³⁰ This paper is explicit on how efficiency emerges as a consequence of sequentially rational redesigns of auction mechanisms.

Appendix A: Proof of Lemma 4

First assume that $\varphi = g \circ h$ is deterministic, i.e., $\varphi(\theta) = g(h(\theta))$ is singleton for all θ . We show that $g(h(\theta)) = \varphi^E[p](\theta)$, for all $\theta \in \Theta$. Denote, for notational simplicity, $Y_i = \{\theta \in \Theta : w(\theta) = i\}$ for all *i*.

Since $\varphi = g \circ h$ is not ex-post dominated by $\varphi^E[p(s, h)]$, $g(h(\theta))$ has to allocate the good to the same buyer as $\varphi^E[p(s, h)](\theta)$ does, for all $s \in h(\theta)$, for all $\theta \in \text{supp}(p)$. Since φ^E satisfies Lemma 1, $\varphi(\theta) = g(h(\theta))$ allocates the good to the same buyer as

^{27.} Milgrom (1987) argues that the core implements the efficient allocation under complete information.

^{28.} Including the English, Vickrey, Dutch, and first-price auctions.

^{29.} The English auction also satisfies a stronger condition of *ex-post implementability*: that one does not want to change one's own behavior even if one knows the behavior of the other players. See for example Bergeman and Morris (2005, 2008).

^{30.} Zheng (2002) is an exception. He characterizes outcome functions that can be implemented with explicit resale markets. See also Haile (2000) for a formal modeling of retrading.

 $\varphi^{E}[p](\theta)$ does, for all $\theta \in \text{supp}(p)$. Thus the partition $\{Y_i\}$ specifies the winner under φ . Since, by EXP-IR, a nonwinner cannot be imposed a strictly positive monetary transfer, the allocation of $g(h(\theta))$ may differ from $\varphi^{E}[p](\theta)$ only in terms of the winner's monetary transfer. Denote the winner's monetary transfer under $g(h(\theta))$ by $m_i(\theta)$. Thus, our task reduces to showing that $m_i(\theta) = m^{E}(i, \theta_{-i}, p)$, for all $\theta \in Y_i$, for all i.

Fix *i*. Since Θ_i is discrete and bounded below, we can order its elements by $\theta_i^0 < \cdots < \theta_i^k < \cdots$. We prove by induction that $m_i(\theta_i^k, \theta_{-i}) = m^E(i, \theta_{-i}, p)$, for all θ_{-i} such that $(\theta_i^k, \theta_{-i}) \in Y_i$, for all $k = 0, 1, \ldots$. Assume that the induction hypothesis holds until the step k - 1, namely

$$m_i(\theta_i^l, \theta_{-i}) = m^E(i, \theta_{-i}, p), \qquad (A.1)$$

for all $(\theta_i^l, \theta_{-i}) \in Y_i \cap \text{supp}(p)$, for all l = 0, ..., k - 1. We show that (A.1) holds also for step k.

Take any $s \in h(\operatorname{supp}(p))$. Since φ does not leave surplus to the winner that could be extracted by $\varphi^{E}(p((x, s), \varphi))$,

$$m_i(\theta_i^k, \theta_{-i}) \ge m^E(i, \theta_{-i}, p(s, h)),$$

for all $(\theta_i^k, \theta_{-i}) \in \operatorname{supp}(p(s, h))$. Since $\operatorname{supp}(p(s, h)) \subseteq \operatorname{supp}(p)$,

$$m^{E}(i, \theta_{-i}, p(s, h)) \ge m^{E}(i, \theta_{-i}, p),$$
(A.2)

for all $(\theta_i^k, \theta_{-i}) \in \text{supp}(p(s, h))$. Noting that (A.2) holds for all $s \in h(\text{supp}(p))$, it follows from the above two conditions that

$$m_i(\theta_i^k, \theta_{-i}) \ge m^E(i, \theta_{-i}, p), \tag{A.3}$$

for all $(\theta_i^k, \theta_{-i}) \in \text{supp}(p)$. It remains to be shown that the weak inequality in (A.3) holds as equality.

By (4),

$$m^{E}(i, \theta_{-i}, p) = \theta_{i}^{k}$$
(A.4)

for all $(\theta_i^k, \theta_{-i}) \in Y_i \cap \text{supp}(p)$ such that $(\theta_i^{k-1}, \theta_{-i}) \notin Y_i$. This has two implications. First,

$$\sum_{\substack{(\theta_i^k,\theta_{-i},)\in Y_i\\(\theta_i^{k-1},\theta_{-i})\in Y_i}} [\theta^k - m^E(i,\theta_{-i},p)] p(\theta_i^k,\theta_{-i})$$

$$= \sum_{\substack{(\theta_i^{k-1},\theta_{-i})\in Y_i\\(\theta_i^k,\theta_{-i},p)]} [\theta^k - m^E(i,\theta_{-i},p)] p(\theta_i^k,\theta_{-i}).$$
(A.5)

Second, by (A.3),

$$m_i(\theta_{-i},\theta_i^k) \ge \theta_i^k$$

for all $m_i(p, \theta_{-i}) = \theta_i^k$ for all $(\theta_i^k, \theta_{-i}) \in Y_i \cap \text{supp}(p)$ such that $(\theta_i^{k-1}, \theta_{-i}) \notin Y_i$. By VETO-IC and this property,

$$\sum_{\substack{(\theta_i^k,\theta_{-i})\in Y_i\\ (\theta_i^k,\theta_{-i})\in Y_i}} \left[\theta_i^k - m_i(\theta_i^k,\theta_{-i})\right] p(\theta_i^k,\theta_{-i})$$

$$\geq \sum_{\substack{(\theta_i^k,\theta_{-i})\in Y_i\\ (\theta_i^{k-1},\theta_{-i})\in Y_i}} \max\left\{\theta_i^k - m_i(\theta_i^{k-1},\theta_{-i}),0\right\} p(\theta_i^k,\theta_{-i}).$$
(A.6)

By the induction hypothesis (A.1),

$$\theta_i^k - m_i(\theta_i^{k-1}, \theta_{-i}) = \theta_i^k - m^E(i, \theta_{-i}, p),$$
(A.7)

for all $(\theta_i^{k-1}, \theta_{-i}) \in Y_i \cap \operatorname{supp}(p)$. By (A.7) and (A.4),

$$\sum_{\substack{(\theta_i^{k-1},\theta_{-i})\in Y_i\cap \text{supp}(p)\\ =\sum_{\substack{(\theta_i^{k-1},\theta_{-i})\in Y_i\cap \text{supp}(p)\\ (\theta_i^{k-1},\theta_{-i})\in Y_i\cap (\theta_i^{k-1},\theta_{-i})\in Y_i}} \left[\theta_i^k - m^E(i,\theta_{-i},p)\right] p(\theta_i^k,\theta_{-i})$$

Thus,

$$\sum_{\substack{(\theta_i^{k-1},\theta_{-i})\in Y_i\\(\theta_i^{k-1},\theta_{-i})\in Y_i}} \max\left\{\theta_i^k - m_i(\theta_i^{k-1},\theta_{-i}), 0\right\} p(\theta_i^k,\theta_{-i})$$

$$\geq \sum_{\substack{(\theta_i^{k-1},\theta_{-i})\in Y_i\\(\theta_i^{k-1},\theta_{-i})\in Y_i}} \left[\theta_i^k - m^E(p,\theta_{-i})\right] p(\theta_i^k,\theta_{-i}).$$

This together with (A.6) implies that

$$\sum_{(\theta_i^k,\theta_{-i})\in Y_i} \left[\theta_i^k - m_i(\theta_i^k,\theta_{-i})\right] p(\theta_i^k,\theta_{-i}) \geq \sum_{(\theta_i^{k-1},\theta_{-i})\in Y_i} \left[\theta_i^k - m^E(i,\theta_{-i},p)\right] p(\theta_i^k,\theta_{-i}).$$

Thus, by (A.5),

$$\sum_{(\theta_i^k,\theta_{-i})\in Y_i} \left[\theta_i^k - m_i(\theta_i^k,\theta_{-i})\right] p(\theta_i^k,\theta_{-i}) \geq \sum_{(\theta_i^k,\theta_{-i})\in Y_i} \left[\theta_i^k - m^E(i,\theta_{-i},p)\right] p(\theta_i^k,\theta_{-i}),$$

and hence

$$\sum_{(\theta_i^k,\theta_{-i})\in Y_i} m_i(\theta_i^k,\theta_{-i}) p(\theta_i^k,\theta_{-i}) \le \sum_{(\theta_i^k,\theta_{-i})\in Y_i} m^E(i,\theta_{-i},p) p(\theta_i^k,\theta_{-i}).$$
(A.8)

Finally, (A.8) implies that (A.3) holds as equality, as desired.

Finally we check the case of random *h*. Note that even if $g \circ h$ is random, it has to allocate the good to the same buyer as $\varphi^{E}[p]$ does. Thus randomness of $g \circ h$ may only concern the monetary transfer *m*. The proof, which is by induction, proceeds along the above lines. It is easy to verify that (A.3) holds also for any randomly generated monetary transfer *m* under $(\theta_{i}^{k}, \theta_{-i})$. Moreover, (A.8) needs to hold for the *expected* transfer $\overline{m} \sum$ under $(\theta_{i}^{k}, \theta_{-i})$. Again, this just means that (A.3) holds as equality for each *m*, thus $m^{E}(i, \theta_{-i}, p)$ is implemented with probability one under all $(\theta_{i}^{k}, \theta_{-i}) \in$ supp(*p*). This completes the proof.

Appendix B: Nonstationarity of σ^{λ}

Let n = 1 and assume the "no gap" case $0 \in \Theta \neq \{0\}$. We construct a seller's choice function σ^{λ} that allows the seller to commit to *any* price $\lambda \in \Theta$. Define a take-it-or-leave-it offer

$$\varphi^{\lambda}(\theta) = \begin{cases} (1,\lambda), & \text{if } \theta \ge \lambda, \\ (0,0), & \text{if } \theta < \lambda. \end{cases}$$

That is, "sell with price λ to any type θ at least λ and not sell to type below λ ". Define choice rule σ^{λ} such that

$$\sigma^{\lambda}[p] = \begin{cases} \varphi^{\lambda}, & \text{if } \{0, \lambda\} \subset \text{supp}(p) \\ \mathbf{1}_{(1,\underline{\theta}(p))}, & \text{otherwise.} \end{cases}$$

That is, if both 0 and λ belong to the support of p, then use φ^{λ} . Otherwise sell with the price equal to the least possible valuation in supp(p).

We claim that σ^{λ} satisfies the one-deviation property:

- (i) If $\{0, \lambda\} \subset \operatorname{supp}(p)$, then $\{(1, \lambda), (0, 0)\} = \sigma^{\lambda}[p]$. Now,
 - $0 < \lambda = \theta(p((1, \lambda), \sigma^{\lambda}[p]))$ and thus $\sigma^{\lambda}[p((1, \lambda), \sigma^{\lambda}[p])] = \mathbf{1}_{(1,\lambda)}$, and
 - $0 = \underline{\theta}(p((0, 0), \sigma^{\lambda}[p]))$ and thus $\sigma^{\lambda}[p((0, 0), \sigma^{\lambda}[p])] = \mathbf{1}_{(0,0)}$.
- (ii) If $\{0, \lambda\} \not\subset \operatorname{supp}(p)$, then $\sigma^{\lambda}(p) = \mathbf{1}_{(1,\theta(p))}$ and $\sigma^{\lambda}(p) = \mathbf{1}_{(0,0)}$, respectively.

Since constant mechanisms do not affect beliefs, σ^{λ} satisfies the one-deviation property. Thus, the seller can commit to any price $\lambda \in \text{supp}(p)$.

We now argue that σ^{λ} does not meet stationarity. To see this, let $\{0, \lambda\} \subset$ supp(*p*). Denote by $\mathbf{1}_0$ the degenerate distribution such that supp($\mathbf{1}_0$) = $\{0\} \not \geq \lambda$. Then supp($\mathbf{1}_0$) \subset supp(*p*). By construction, $\sigma^{\lambda}[p] = \varphi^{\lambda}$ and $\sigma^{\lambda}[\mathbf{1}_0] = \mathbf{1}_{(1,0)}$. Therefore, $\sigma^{\lambda}[p](0) = (0, 0)$ and $\sigma^{\lambda}[\mathbf{1}_0](0) = (1, 0)$. However, v((1, 0)) = v((0, 0)) = 0, violating stationarity.

This result is analogous to Ausubel and Deneckere (1989), who show that any price can be supported in equilibrium in the one-sided offers bargaining game when the discount factor δ tends to 1. Strategies needed for these equilibria are complicated—that is, nonstationary.

References

- Adams, Christopher (2005). "Roundtable on the Economics of Internet Auctions: An Executive Summary." Federal Trade Comission Bureau of Economics.
- Ausubel, Larry and Peter Cramton (1999). "The Optimality of Being Efficient." Working paper, University of Maryland.
- Ausubel, Larry and Raymond Deneckere (1989a). "Reputation in Bargaining and Durable Goods Monopoly." *Econometrica*, 57, 511–532.
- Ausubel, Larry and Raymond Deneckere (1989b). "A Direct Mechanism Characterization of Sequential Bargaining with One-Sided Incomplete Information." *Journal of Economic Theory*, 48, 18–46.
- Ausubel, Larry, Peter Cramton and Raymond Deneckere (2002). "Bargaining with Incomplete Information." In *Handbook of Game Theory, Vol. 3*, edited by Robert Aumann and Sergiu Hart. Elsevier, chap. 50.
- Bergemann, Dirk and Stephen Morris (2005). "Robust Mechanism Design." *Econometrica*, 73, 1771–1813.
- Bergemann, Dirk and Stephen Morris (2008). "Ex Post Implementation." *Games and Economic Behavior*, 63, 527–566.
- Bester, Helmut and Roland Strausz (2000). "Imperfect Commitment and the Revelation Principle: The Multi-agent Case." *Economics Letters*, 69, 165–171.
- Bester, Helmut and Roland Strausz (2001). "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case." *Econometrica*, 69, 1077–1098.
- Boyd, Colin and Wenbo Mao (2000). "Security Issues for Electronic Auctions." Hewlett Packard TR-HPL-2000.
- Calzolari, Giacomo and Alessandro Pavan (2006a). "Monopoly with Resale." *RAND Journal of Economics*, 37, 362–375.
- Calzolari, Giacomo and Alessandro Pavan (2006b). "On the Optimality of Privacy in Sequential Contracting." *Journal of Economic Theory*, 130, 168–204.
- Chakraborty, Indranil and Georgia Kosmopoulou (2004). "Auctions with Shill Bidding." *Economic Theory*, 24, 271–287.
- Crémer, Jacques and Richard P. McLean (1988). "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica*, 56, 6, 1247–1258.
- Coase, Ronald (1972). "Durability and Monopoly." Journal of Law and Economics, 15, 143-149.
- Forges, Francoise (1993). "Some Thoughts on Efficiency and Information." In *Frontiers of Game Theory*, edited by Ken Binmore, Alan Kirman and Piero Tani. MIT Press, pp. 133–153.
- Forges, Francoise (1998). "Ex Post Individually Rational Trading Mechanisms." Working paper, THEMA, University of Cergy-Pontoise.
- Fudenberg, Drew, David Levine, and Jean Tirole (1985). "Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information." In *Game Theoretic Models of Bargaining*, edited by Alvin Roth. Cambridge University Press, pp. 73–98.
- Garratt, Rod and Thomas Troger (2006). "Speculation in Standard Auctions with Resale." *Econometrica*, 74, 753–769.
- Green, Jerry and Jean-Jacques Laffont (1987). "Posterior Implementability in a Two-Player Decision Problem." *Econometrica*, 55, 69–94.
- Gresik, Thomas (1991). "Ex Ante Efficient, Ex Post Individually Rational Trade." *Journal of Economic Theory*, 53, 131–145.
- Gresik, Thomas (1996). "Incentive Efficient Equilibria of Two-Party Sealed-Bid Bargaining Games." Journal of Economic Theory, 68, 26–48.
- Gul, Faruk, Hugo Sonnenschein, and Robert Wilson (1986). "Foundations of Dynamic Monopoly and The Coase Conjecture." *Journal of Economic Theory*, 39, 155–190.
- Hafalir, Isa and Vijay Krishna (2008). "Asymmetric Auctions with Resale." *American Economic Review*, 98(1), 87–112.
- Izmalkov, Sergei (2005). "Shill Bidding and Optimal Auctions." Working paper, Raymond MIT.

- Klemperer, Paul (1999). "Auction Theory: a Guide to the Literature." *Journal of Economic Surveys*, 13, 237–286.
- Krishna, Vijay R. (2007). "Communication in Games of Incomplete Information: The Two-Player Case." *Journal of Economic Theory*, 132, 584–592.
- Krishna, Vijay and Motty Perry (1998). "Efficient Mechanism Design." Working paper, Penn State University.
- Lopomo, Giuseppe (1998). "The English Auction is Optimal Among Simple Sequential Auctions." Journal of Economic Theory, 82, 144–166.
- Lopomo, Giuseppe (2001). "Optimality and Robustness of the English Auction." *Games and Economic Behavior*, 36, 219–245.
- Lucking-Reiley, David (2000a). "Auctions on the Internet: What's Being Auctioned, and How." *Journal of Industrial Economics*, 48, 227–252.
- Lucking-Reiley, David (2000b). "Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-Commerce." *Journal of Economic Perspectives*, 14, 183–192.
- Matthews, Steven (1987). "Comparing Auctions for Risk Averse Buyers: a Buyer's Point of View." *Econometrica*, 89, 633–646.
- Matthews, Steven and Andrew Postlewaite (1989). "Pre-play Communication in Two-Person Sealed-Bid Double Auctions." *Journal of Economic Theory*, 89, 238–263.
- McAdams, David and Michael Schwarz (2007a). "Credible Sales Mechanisms and Intermediaries." American Economic Review, 97, 260–276.
- McAdams, David and Michael Schwarz (2007b). "Who Pays When Auction Rules are Bent?" International Journal of Industrial Organization, 25, 1144–1157.
- McAfee, Preston and Phil Reny (1992). "Correlated Information and Mechanism Design." *Econometrica*, 60, 395–421.
- McAfee, Preston and Daniel Vincent (1997). "Sequentially Optimal Auctions." *Games and Economic Behavior*, 18, 246–276.
- Milgrom, Paul (1987). "Auction Theory." In *Advances in Economic Theory: Fifth World Congress*, edited by Truman Bewley. Cambridge University Press, pp. 1–32.
- Milgrom, Paul and Robert Weber (1982). "A Theory of Auctions and Competitive Bidding." *Econometrica* 50, 1089–1122.
- Myerson, Roger (1981). "Optimal Auction Design." Mathematics of Operation Research, 6, 58-73.
- Myerson, Roger and Mark Satterthwaite (1983). "Efficient Mechanisms for Bilateral Trading." Journal of Economic Theory, 29, 265–281.
- Ockenfels, Axel, David Reiley, and Sadrieh Abdolkarim (2007). "Online Auctions." NBER Working Paper No. 12785.
- Rothkopf, Michael, Thomas Teisberg, and Edward Kahn (1990). "Why Are Vickrey Auctions Rare?" Journal of Political Economy 98, 94–109.
- Skreta, Vasiliki (2006). "Sequentially Optimal Mechanisms." *Review of Economic Studies*, 73, 1085–1111.
- Skreta, Vasiliki (2010). "Optimal Auction Design under Non-commitment." Working paper, NYU. Zheng, Charles (2002). "Optimal Auction with Resale." *Econometrica*, 70, 2197–2224.