Optimal Search Auctions[1]

Jacques Crémer[2] Yossi Spiegel[3] Charles Z. Zheng[4]

First draft: September 24, 2002

This draft: December 28, 2005

[1]This version supersedes a draft titled "Optimal Auctions from a Search-Theoretic Viewpoint." We are grateful for comments received from Eddie Dekel, Rich McLean, Andrew McLennan, Josh Pinkston, Michael Schwarz, Asher Wolinsky, Lixin Ye, an associate editor, and two referees. Yossi Spiegel gratefully acknowledges the financial support of the IIBR and Charles Zheng thanks the NSF for grant SES-0214471.

[2]IDEI-GREMAQ, Université de Toulouse, Tel: +33 (0)5 61 12 86 34, email: jacques@cremeronline.com

[3]Recanati Graduate School of Business Administration, Tel Aviv University, Ramat Aviv, Tel Aviv, 69978, Israel, Tel: 972-3-640-9063, email: spiegel@post.tau.ac.il, http://www.tau.ac.il/~spiegel

[4]Department of Economics, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208, Tel: 847-491-8213, email: czheng@northwestern.edu.

Abstract

We study the design of profit maximizing single unit auctions under the assumption that the seller needs to incur costs to contact prospective bidders and inform them about the auction. With independent bidders' types and possibly interdependent valuations, the seller's problem can be reduced to a search problem in which the surplus is measured in terms of virtual utilities minus search costs. Compared to the socially efficient mechanism, the optimal mechanism features fewer participants, longer search conditional on the same set of participants, and inefficient sequence of entry.

1 Introduction

Almost all the auction literature assumes that the set of bidders is either exogenous or determined in advance before the auction begins. However, auctions based on this assumption are in general suboptimal if the seller incurs costs when contacting prospective bidders. In this paper we study profit maximizing auctions in the presence of these costs. We characterize the order in which bidders are approached, and study the inefficiencies that arise due to the sequential nature of the process and due to the bidders' private information about their values.

We consider a seller of a single indivisible good who faces a finite set of bidders. The bidders' types are independently drawn, with ex post valuations interdependent across bidders as in Myerson [11].[1] Initially, prospective bidders are not even aware of the seller's intention to sell the good. To attract their attention and allow them to participate, the seller must contact them and provide them with all the necessary information—in Section 2 we discuss several interpretations of this assumption. After being contacted by the seller and informed about the good for sale, each bidder privately learns his type before deciding whether or not to participate in the seller's mechanism. Given that contacting prospective bidders is costly, it is generally not optimal to contact all bidders at once. For instance, if the expected valuation of an early bidder turns out to be sufficiently high, it is best to end the mechanism immediately and sell him the good without incuring further costs. Hence the seller designs a *search mechanism* that, contingent on history, specifies the order in which prospective bidders are contacted, the time at which the process ends, and the payments made by the participating bidders.

In Section 2, we introduce the model and notations for search mechanisms. In Section 3.1, we prove one of the main results, Theorem 1: the seller's problem can be reduced to a standard search problem in which the payoff from search is measured in terms of the winner's' virtual utility rather than his actual utility. This result is nontrivial despite its well known counterpart in the static framework; the main complication in the proof is that the bidders' incentive constraints depend on the dynamic stochastic nature of the seller's optimal search problem.

The optimal search mechanisms we study here give rise to new types of distortions that are completely absent in static mechanism design problems. In Sections 4.1–4.3, we present three of these distortions: we show that asymmetry of information leads to fewer participants, longer

[1] Our model can be interpreted as a procurement model where a buyer wishes to buy a good from one of several suppliers whose types are independently drawn from bidder-specific distributions, and whose costs are interdependent.

search conditional on the same set of participants, and inefficient sequence of entry.

Section 4.4 presents another feature of optimal search auctions: When the ex post values of bidders are interdependent, the seller wants to delay the participation of "influential" bidders, whose types have a strong effect on the willingness of others to pay for the good.

Our findings extend existing results in traditional mechanism design theory by endogenizing the set of participants through a stochastic, history-contingent, search procedure.[2] Myerson [11] has characterized the optimal (profit maximizing) auctions in the case where the seller incurs no search cost when contacting potential bidders, so there is no loss to assume that they all participate. Hence his solution arises as a special case in our framework. McAfee and McMillan [9] characterize optimal search mechanisms in the special case where bidders are ex ante symmetric in terms of the search cost and the distribution of their types. Hence the sequence of entry and the resulting distortions which play a major role in our paper do not matter in their model. We allow bidders to be ex ante asymmetric so the sequence of entry is important, and we find some features of the optimal sequence in Sections 3.2, 4.3, and 4.4. Moreover, McAfee and McMillan assume that there is no discounting so that it does not hurt the seller to contact only one bidder at a time. We allow discounting so the seller may need to contact a group of bidders in a period.

Two other papers have also applied search theory to auction design. Burguet [4] considers a procurement model with private-value ex ante symmetric bidders who must decide whether or not to participate before knowing their types. In Crémer, Spiegel, and Zheng [6], we generalize Burguet's results, in the context of an auction model, by allowing for the entry of multiple bidders at each stage and allowing for both interdependent valuation and asymmetric bidders. Unlike these two papers, which assume ex ante participation constraints, this paper uses interim participation constraints, as bidders are privately informed about their types during their participation decisions.

To the best of our knowledge, other than McAfee and McMillan [9] and Burguet [4], we are the only ones optimizing over the rules of the auctions (as opposed to comparing specific auction formats) given the constraints of costly participation. Other important works on mechanism design with information acquisition or costly participation include Bergemann and Pesendorfer [1], Bergemann and Välimäki [2], Levin and Smith [8], and Ye [18]. (See Bergemann and Välimäki [3]

[2]Our analysis may also contribute to the optimal search literature by highlighting a new parallel search problem where the ex post social surplus from selling the good to a bidder depends on the signals of other potential bidders, whether they have participated or not. We have not found any work in that literature that considers this problem. Weitzman [16] and Vishwanath [15] considered only private values. Even for private values, the literature has no general characterization of optimal search procedures that allow multiple entrants per period, which we consider.

for a recent survey.) Unlike these papers, which assume that the agents' participation decisions are made independently of each other, this paper allows an agent's entry decision to depend on the history of a mechanism. With this dynamic feature, our paper is somewhat related to Compte and Jehiel [5] and Rezende [13], who analyze the effect of information acquisition conducted during ascending price auctions.[3]

2 The Model

2.1 Search costs

A seller wants to sell an indivisible good to one out of a finite set I of prospective bidders. Initially, bidders do not know the seller's intention to sell the good and are not aware of the auction setting (the rules of the auction, the number and identity of other bidders, and the distribution of bidders' valuations). To bring this information to a bidder i's attention, the seller incurs a bidder-specific fixed cost $c_i > 0$, which we call *search cost*. While learning the information, the bidder also privately learns his own type which affects (but is not necessarily equal to) his ex post valuation.

The cost c_i has several possible interpretations. First, if the seller's good is very complex (e.g., the controlling block of a state-owned enterprise), the seller may need to meet potential bidders in person and describe the good in detail.[4] Second, although we consider an auction environment, our framework can be easily modified to a procurement environment in which a procurer wishes to procure an indivisible good from a set of potential suppliers. If the procurer's

[3] There are a few other interesting papers on specific auction formats with information acquisition or costly participation, including Gal, Landsberger, and Nemirovski [7], Stegeman [14], Pesendorfer and Wolinsky [12], and Wolinsky [17]. However, unlike in our paper, the participation or information acquisition decisions in these papers are not coordinated by a principal.

[4] If the seller has goals other than profit maximization, he may also have to meet with potential bidders in person in order to ensure that they meet certain criteria (e.g., ensure that a privatized state-owned enterprise will be controlled by a qualified manager). In this paper however we do not consider this possibility explicitly; doing so may require us to consider additional dimensions of the buyers' private information beyond their valuation of the seller's good (e.g., their level of "competence"). Moreover, in many examples it is likely that bidders may also have to incur costs in order to participate in the seller's mechanism. In this paper however, we focus exclusively on the case where only the seller needs to incur search costs. In Crémer, Spiegel and Zheng [6], we consider the other extreme case and assume that the buyers need to incur costs to learn their types before they can participate in the seller's mechanism. It would be interesting in future research to study the harder case where both the seller and the buyers need to incur search costs.

needs are complex and hard to describe, he would need to understand exactly what each supplier can offer. For instance, consider a firm that wants to outsource a custom-made component; in some cases, rather than sending the description to a prospective supplier and asking for a price quote, it could be more efficient to ask for a description of the supplier's manufacturing facilities and explain what type of steps need to be taken with these specific facilities in order to produce the good. The supplier can then provide a quote.[5]

2.2 Utility functions and types

The value of the good to the seller is x_0. For each bidder i, nature draws a *type* x_i from a commonly known distribution F_i, with density f_i and support $X_i = [\underline{x}_i, \overline{x}_i]$, with $f_i > 0$ over the interior of X_i. Types are independent across i. Denote

$$x := (x_i)_{i \in I} \in \times_{i \in I} X_i =: X.$$

As in Myerson [11], given any x, bidder i's value of the good is equal to

$$u_i(x) := x_i + \sum_{j \in I \setminus i} e_{ij}(x_j),$$

where e_{ij} is a commonly known real function that reflects bidder j's influence on bidder i's value. Hence, bidder i's value for the good depends not only on his own type, x_i, but also on the types of other bidders through the functions $\{e_{ij}(\cdot)\}_{j \in I \setminus i}$. Everyone's discount factor is $\delta \in (0, 1]$. If bidder i pays p_i^t dollars in period t, then his utility from the viewpoint of period s is $\delta^{t'-s} u_i(x) - \sum_{t=s}^{\infty} \delta^{t-s} p_i^t$ if he gets the good in period $t' \geq s$, and $-\sum_{t=s}^{\infty} \delta^{t-s} p_i^t$ if he never gets it. The seller uses the same discount factor to evaluate his present discounted profit.

2.3 Search mechanisms

When the seller needs to incur costs to contact specific bidders, it is in general suboptimal (both socially and from the seller's viewpoint) to commit in advance to a fixed set of participants without knowing the bidding history. Hence the seller picks a contingent plan that, based on the messages of the "incumbents", specifies whether the seller should (i) stop the mechanism and either keep the

[5]Another example might be a movie producer looking for a location to shoot a new movie. In order to get price quotes from the various potential locations, the producer needs to examine the exact facilities that each location can offer. Only then can the producer tell which facilities it would need and obtain a price quote.

good or allocate it to an incumbent bidder, or (ii) continue and invite new bidders. Coupled with a payment scheme, such a contingent plan is called a *search mechanism*. Note that parallel search is allowed since the seller can invite several entrants at once.

At the start of period 1, the seller contacts a set of entrants. If an entrant agrees to participate, he signs with the seller a binding contingency contract. Since he is privately informed before signing the contract, a bidder's participation constraint is interim.[6] Each period-1 entrant then sends a message. Given these messages, either the mechanism stops and the seller keeps the good or allocates it to a period-1 entrant, or the mechanism continues to period 2 and more entrants are invited. Depending on the information disclosure policy that the seller adopts as part of the mechanism, each period-2 entrant is told none, or part of, or all of the messages sent by previous entrants. Therefore, a new entrant need not even know what the current period is. Given the disclosed information, each period-2 entrant decides whether to participate; if he does, he sends a message. Depending on the messages sent in periods 1 and 2, the mechanism either continues and a new set of entrants is invited to participate in period 3, or the mechanism stops. In the latter case, the seller keeps the good or allocates it to a period-1 or period-2 entrant. The mechanism continues in a similar fashion until it stops and the good is allocated.

2.4 Revelation search mechanisms

A *revelation search mechanism* is a search mechanism in which each bidder i's message space is i's type space. A search mechanism, as a multistage game, is *equilibrium feasible* if it has a perfect Bayesian equilibrium (PBE). A revelation search mechanism is *incentive feasible* if it has a PBE where every invited bidder participates and is truthful.

The next lemma allows us to restrict attention to revelation search mechanisms without loss. Its proof is similar to the proof of the standard revelation principle and hence is omitted.

Lemma 1 (Revelation Principle for Search Mechanisms). *There exists an incentive feasible revelation search mechanism that replicates the equilibrium outcome of any equilibrium feasible search mechanism.*

Given a revelation search mechanism, suppose that the profile of realized values is $x \in X$ and every invited bidder participates and is truthful. Then the mechanism induces the following

[6] We assume that a seller cannot make a buyer commit to a payment plan before the buyer knows his type. Otherwise, the buyer would be besieged by dishonest sellers selling him fake projects or goods of no value whatsoever.

objects. (The formal definition is in the appendix.)

$$\begin{aligned}
E^t(x) &:= \text{the set of potential bidders who enter the mechanism in period } t;\\
q_i(x) &:= \text{the probability with which player } i \text{ (bidder or seller) consumes the good;}\\
p_i(x) &:= \text{the total payment made by bidder } i \text{ discounted back to the period}\\
&\quad\;\; \text{at which } i \text{ enters the mechanism;}\\
H_i(x) &:= \text{the event (set of possible states) that bidder } i \text{ knows has occurred}\\
&\quad\;\; \text{when } i \text{ is entering the mechanism;}\\
t_i(x) &:= \text{the period at which player } i \text{ enters the mechanism } (i \in E^{t_i(x)}(x));\\
\tau(x) &:= \text{the period at which the search terminates;}\\
I^t(x) &:= \cup_{s=1}^t E^s(x) \;=\; \text{the set of incumbents at the end of period } t.
\end{aligned}$$

A revelation search mechanism is denoted $((E^t)_{t=1}^\infty, q, (p_i, H_i)_{i\in I})$. A *search procedure* is the operation-research part of a revelation search mechanism; it determines the set of entrants in each period and the identity of the winner of the good but not how much to charge or what information to disclose. Denote a search procedure by $((E^t)_{t=1}^\infty, q)$.[7]

The sequential nature of a search mechanism is formulated into the following constraints (which are implied by the formal definition in the appendix):

1. E^1 is constant on X, i.e., the set of entrants in period 1 is determined without any message.

2. If two realized profiles x and x' of bidder-values generate the same history up to period $t \geq 1$, i.e, $E^s(x) = E^s(x')$ for all $s = 1, \ldots, t$, and if $x_i = x_i'$ for all incumbents $i \in I^t(x)$, then x and x' induce the same decisions for period $t+1$:

   a. The set of entrants in period $t+1$ is the same and this set is a subset of potential bidders who have not yet entered, i.e., $E^{t+1}(x) = E^{t+1}(x') \subseteq I \setminus I^t(x)$.

   b. If $t = \tau(x)$, then $q(x) = q(x')$ and $p_i(x) = p_i(x')$ for every potential bidder i; i.e., if the mechanism stops at period t, then the allocation rule is the same for x and x', and the payment is the same for x and x'.

   c. For any i who enters the mechanism in period $t+1$, $H_i(x) = H_i(x')$, i.e., the news disclosed to entrant i is the same for x and x'.[8]

[7]Note: $((E^t)_{t=1}^\infty)$ determines the period $\tau(x)$ in which the search ends: $\tau(x) = \max\{s = 1, 2, \ldots : E^s(x) \neq \varnothing\}$.

[8]Constraint 2.c implies that a participant cannot learn about the types of those who have not participated. To prove

3. The good cannot go to bidders who do not participate in the mechanism, nor can the seller collect payments from such bidders, i.e., $i \notin I^{\tau(x)}(x) \Rightarrow q_i(x) = p_i(x) = 0$.

The functions $(H_i)_{i \in I}$ constitute the disclosure policy of the mechanism. In a *non-disclosure* policy, a bidder upon entry learns only the fact that the mechanism is contacting him at the current period without knowing what the current period is. In a *full disclosure* policy, for any profile x of realized values and in any period t, every entrant knows the sequence $(E^s(x))_{s=1}^t$ of entry up to now, as well as the reported type x_j of every incumbent $j \in I^{t-1}(x)$.

Whatever the disclosure policy, the bidders are told the mechanism. A bidder makes the participation decision only after he has been informed of his type and whatever the disclosure policy reveals to him, including all the rules of the mechanism.

The above notations implicitly assume that the seller does not fully randomize on the search procedure. This assumption causes no loss of generality, because the seller and bidders are all risk-neutral. Our formal definition in the Appendix does allow full randomization.

2.5 Notions of optimal search mechanisms

Given any search procedure $((E^t)_{t=1}^\infty, (q_i)_{i \in I})$, if all invited bidders participate and are truthful, and if the seller gets the entire surplus, then the seller's expected profit discounted to period 1 is

$$\Pi((E^t)_{t=1}^\infty, (q_i)_{i \in I} \mid (u_i)_{i \in I}) := \mathbb{E}_x \left[\delta^{\tau(x)-1} \left[\sum_{i \in I} q_i(x) \left(u_i(x) - x_0 \right) \right] - \sum_{t=1}^{\infty} \delta^{t-1} \sum_{i \in E^t(x)} c_i \right]. \quad (1)$$

Note that there is no need to quantify the first sum by the restriction that i is a participant, because if i is not a participant then $q_i(x) = 0$ by constraint 3 in the previous subsection.

In traditional search theory there is no asymmetric information once the search cost has been incurred. Thus, optimal search amounts to maximizing $\Pi((E^t)_{t=1}^\infty, (q_i)_{i \in I} \mid (u_i)_{i \in I})$ over all search procedures. We call this unconstrained maximization problem *symmetric-information search problem relative to payoffs* $(u_i)_{i \in I}$ and say that its solution is *symmetric-information optimal relative to payoffs* u_i. In our auction environment, by contrast, after the seller incurs a search cost and contacts a bidder, the bidder becomes privately informed about his type. Hence the seller needs to design a search mechanism that induces the bidders to reveal their private information truthfully.

that, suppose at state x, i enters in period $t+1$ and j has not entered by the end of period $t+1$ (i.e., $i \in E^{t+1}(x)$ and $j \notin I^{t+1}(x)$). For any possible type x'_j of j, denote (x_{-j}, x'_j) for the state such that the type of j is x'_j and the type of everyone else is the same in x. By 2.c, $H_i(x) = H_i(x_{-j}, x'_j)$, hence i has no way to update about j's type.

A search procedure $((E^t)_{t=1}^{\infty}, (q_i)_{i\in I})$ is *optimal* (profit maximizing) if $\Pi((E^t)_{t=1}^{\infty}, (q_i)_{i\in I} \mid (u_i)_{i\in I})$ is maximized over all search procedures subject to interim participation and incentive compatibility constraints. Similarly, a search mechanism is *optimal* if it implements an optimal search procedure.

3 Optimal search mechanisms

For every potential bidder $i \in I$ and any possible realized profile $x := (x_i)_{i\in I}$ of bidder-values, define the (ex post) virtual utility of bidder i to be

$$V_i(x) := x_i - \frac{1 - F_i(x_i)}{f_i(x_i)} + \sum_{j\in I\setminus i} e_{ij}(x_j). \tag{2}$$

Following most of the optimal auction literature, we make the following assumption which extends the usual monotone hazard rate assumption to the case of interdependent values.

Assumption 1. For any potential bidders i and j, $x_i - \frac{1-F_i(x_i)}{f_i(x_i)}$ and $e_{ji}(x_i)$ are differentiable functions of x_i on X_i, their derivatives are uniformly bounded, and $\frac{d}{dx_i}\left(x_i - \frac{1-F_i(x_i)}{f_i(x_i)}\right) > e'_{ji}(x_i) \geq 0$ over the interior of X_i.

Theorem 1. *If Assumption 1 holds, then (a) disclosure policies do not affect the seller's expected profit, and (b) there is an optimal search mechanism that uses the symmetric-information optimal search procedure relative to virtual utility functions* $(V_i)_{i\in I}$.

Hence a profit-maximizing seller just needs to solve a *distorted symmetric-information search problem*, where real utilities $(u_i)_{i\in I}$ are replaced by virtual utilities $(V_i)_{i\in I}$. Moreover, the seller can pick any disclosure policy ranging from non-disclosure to full disclosure. Once he finds a search procedure that solves the distorted problem and arbitrarily picks a disclosure policy, the seller just needs to implement them with a payment scheme satisfying the familiar envelope formula.

The irrelevance of disclosure policies might be unexpected. Here we sketch the economic reasoning behind the proof. Note that no matter how complicated a search procedure is, from the seller's and each bidder's viewpoint, what matters is the bidder's discounted expected probability of winning, conditional on the bidder's own type and the bidder's information set about the types of the other bidders. The latter is disclosed to the bidder by the revelation search mechanism only when the bidder is about to act (who acts only once in the mechanism). As long as an appropriate monotonicity condition holds, the seller can induce each bidder to be truthful by offering the bidder a payment plan conditional on the bidder's information set. Hence the seller's discounted surplus extracted from each bidder i is uniquely determined by the bidder's type, x_i, and the bidder's information set, $\mathbb{H}_i$. Before the search procedure starts, the seller's discounted expected surplus extracted from bidder i is calculated by integrating the surplus extracted from i across all possible information sets $\mathbb{H}_i$ that the mechanism may disclose to this bidder. Note that these information sets constitute a partition on the set of possible realized value-profiles x_{-i} of the other bidders. Hence this integration is the same as integrating across all the possible realized values of x_{-i}, which is independent of the disclosure policy. Thus, the surplus that the seller extracts from each bidder i is uniquely determined by the bidder's type x_i and is independent of the disclosure policy.

3.1 The proof of the theorem

We prove Theorem 1 in three steps. First, in Subsection 3.1.1, we show that the seller's optimal discounted expected profit is bounded from above by the optimal payoff from the distorted symmetric-information search problem. Second, we show in Subsection 3.1.2 that the solution for this distorted symmetric-information problem is incentive compatible given full disclosure policy. Finally, based on step two, we show in Subsection 3.1.3 that the seller can achieve the upper bound in step one via full disclosure. Since other disclosure policies cannot yield less expected profit than full disclosure, any disclosure policy, coupled with a solution for the distorted symmetric-information search problem, is optimal for the seller.

3.1.1 Step one: Finding necessary conditions for incentive feasibility and optimality

By Lemma 1, we can confine attention to revelation search mechanisms. Consider a revelation search mechanism, described by the notations E^t, q_i, H_i, t_i, and τ introduced in Section 2.4.

Assume that bidder i is invited at some period. Let $\mathbb{H}_i$ be the information set disclosed to bidder i before he reports his type. If he reports that his type is $\hat{x}_i$, then from the viewpoint of the current period (which depends on the realized state and may be unknown to i), the discounted expected value of his winning probability is

$$Q_i(\hat{x}_i \mid \mathbb{H}_i) = \mathbb{E}_{x_{-i}}\left[\delta^{\tau(\hat{x}_i, x_{-i}) - t_i(\hat{x}_i, x_{-i})} q_i(\hat{x}_i, x_{-i}) \mid \mathbb{H}_i = H_i(\hat{x}_i, x_{-i})\right], \tag{3}$$

and the discounted expected value of other bidders' influence on bidder i's utility is

$$e_{-i}(\hat{x}_i \mid \mathbb{H}_i) = \mathbb{E}_{x_{-i}}\left[\delta^{\tau(\hat{x}, x_{-i}) - t_i(\hat{x}, x_{-i})} q_i(\hat{x}, x_{-i}) \sum_{j \in I \setminus i} e_{ij}(x_j) \middle| \mathbb{H}_i = H_i(\hat{x}_i, x_{-i})\right]. \tag{4}$$

The discounted expected value of bidder i's payment from this viewpoint is calculated analogously and denoted by $P_i(\hat{x}_i \mid \mathbb{H}_i)$. If bidder i's realized type is x_i, then his discounted expected utility from the viewpoint of the current period is

$$u_i(\hat{x}_i \mid x_i, \mathbb{H}_i) = x_i Q_i(\hat{x}_i \mid \mathbb{H}_i) + e_{-i}(\hat{x}_i \mid \mathbb{H}_i) - P_i(\hat{x}_i \mid \mathbb{H}_i). \tag{5}$$

Given the independence of bidders' types, none of $Q_i(\hat{x}_i \mid \mathbb{H}_i)$, $P_i(\hat{x}_i \mid \mathbb{H}_i)$ and $e_{-i}(\hat{x}_i \mid \mathbb{H}_i)$ vary with bidder i's actual type. Thus, each bidder i's objective function takes the quasilinear form $x_i A_i(\hat{x}_i) + B_i(\hat{x}_i)$, standard in auction theory. This quasilinear form, coupled with standard techniques (e.g., Myerson [11, Lemma 2]), yields the next lemma.

Lemma 2. *The seller's problem is equivalent to maximizing his discounted expected profit among all pairs of search procedures $((E^t)_{t=1}^{\infty}, (q_i)_{i\in I})$ and disclosure rules $(H_i)_{i\in I}$ subject to the following constraints for any $i \in I$, any $x_i \in X_i$, and any $\mathbb{H}_i$ in the range of H_i ($\{H_i(x) : x \in X\}$):*

$$\begin{aligned} \text{the function } Q_i(\cdot \mid \mathbb{H}_i) \quad &\text{is} \quad \text{nondecreasing}; && (6)\\ u_i(x_i \mid x_i, \mathbb{H}_i) \quad &= \quad u_i(\underline{x}_i \mid \underline{x}_i, \mathbb{H}_i) + \int_{\underline{x}_i}^{x_i} Q_i(z \mid \mathbb{H}_i)dz; && (7)\\ u_i(\underline{x}_i \mid \underline{x}_i, \mathbb{H}_i) \quad &= \quad 0. && (8) \end{aligned}$$

A revelation search mechanism is incentive feasible if (6) is satisfied and the payment scheme satisfies the envelope formula,

$$P_i(x_i \mid \mathbb{H}_i) = x_i Q_i(x_i \mid \mathbb{H}_i) + e_{-i}(x_i \mid \mathbb{H}_i) - \int_{\underline{x}_i}^{x_i} Q_i(z \mid \mathbb{H}_i)dz \tag{9}$$

for any $i \in I$, any $x_i \in X_i$, and any $\mathbb{H}_i$ in the range of H_i.

The next lemma is similar to the "integration by parts" routine in optimal auction theory.

Lemma 3. *If a revelation search mechanism with any disclosure policy is incentive feasible, then the seller's discounted expected profit is $\Pi((E^t)_{t=1}^{\infty}, (q_i)_{i\in I} \mid (V_i)_{i\in I})$, where $((E^t)_{t=1}^{\infty}, (q_i)_{i\in I})$ is the search procedure of the mechanism.*

Proof: Let bidder i enter at period t and be informed of information set $\mathbb{H}_i$. By Eqs. (5), (7), and (8), the seller's expected net profit extracted from bidder i discounted to period t is

$$\mathbb{E}_{x_i}\left[(x_i - x_0)Q_i(x_i \mid \mathbb{H}_i) + e_{-i}(x_i \mid \mathbb{H}_i) - \int_{\underline{x}_i}^{x_i} Q_i(z \mid \mathbb{H}_i)dz - c_i\right].$$

By a standard argument (e.g., Myerson [11, Lemma 3]), this is equal to

$$\mathbb{E}_x \left[\delta^{\tau(x)-t} q_i(x) \left(V_i(x) - x_0\right) - c_i \mid \mathbb{H}_i = H_i(x)\right],$$

where we used Eqs. (2), (3), and (4). Viewed from period 1, the period t at which bidder i enters the mechanism is a random variable uniquely determined by the profile x of realized values: $t = t_i(x)$. Thus, viewed from period 1, the seller's expected profit extracted from bidder i is

$$\begin{aligned} &\mathbb{E}_t \left[\mathbb{E}_{\mathbb{H}_i} \left[\mathbb{E}_x \left[\delta^{\tau(x)-1} q_i(x) \left(V_i(x) - x_0\right) - \delta^{t-1} c_i\right] \mid \mathbb{H}_i = H_i(x)\right] \mid t = t_i(x)\right] \\ &\qquad = \mathbb{E}_x \left[\delta^{\tau(x)-1} q_i(x) \left(V_i(x) - x_0\right) - \sum_{t:\ t = t_i(x)} \delta^{t-1} c_i\right]. \quad (10) \end{aligned}$$

Summing the right-hand side (10) over all $i \in I$ (and noting that $t = t_i(x) \Leftrightarrow i \in E^t(x)$), we get

$$\mathbb{E}_x \left[\delta^{\tau(x)-1} \left[\sum_{i \in I} q_i(x) \left(V_i(x) - x_0\right)\right] - \sum_{t=1}^{\infty} \delta^{t-1} \sum_{i \in E^t(x)} c_i\right], \quad (11)$$

which is equal to $\Pi((E^t)_{t=1}^{\infty}, (q_i)_{i \in I} \mid (V_i)_{i \in I})$, defined in Eq. (1), with V_i replacing u_i. ■

Remark: We would have obtained the traditional recipe of optimal auction by now had search costs been zero: for every state x, set $q_i(x) := 1$ for the bidder i whose virtual utility is highest among all bidders and exceeds x_0. This, however, is in general infeasible for a search mechanism, because the seller does not know the realized types of bidders who have not yet been contacted.

3.1.2 Step two: Verifying the incentive feasibility condition

The purpose of the second step is to show that the solution for the seller's distorted symmetric-information problem is incentive compatible given full disclosure policy. To this end, recall that, in the distorted symmetric-information search problem, the seller tries to maximize the discounted expected value of virtual utility of the winner of the good minus search costs. Denote:

$((\tilde{E}^t)_{t=1}^{\infty}, (\tilde{q}_i)_{i \in I})$:= a search procedure that solves the distorted symmetric-information search problem.

Next, at any period, denote:

$$\begin{aligned}
J &:= \text{the set of incumbents, bidders who have entered;} \\
x_J &:= (x_i)_{i\in J}; \quad x_{-J} := (x_i)_{i\notin J}; \\
\pi(J, x_J) &:= \text{the seller's optimal discounted expected payoff} \\
&\quad\ \text{given the state variable } (J, x_J); \\
\pi_+(J, x_J) &:= \text{the seller's optimal discounted expected payoff} \\
&\quad\ \text{from continuing the search given } (J, x_J) \\
&= \delta \max_{K\subseteq I\setminus J} \left[\mathbb{E}_{x_K} \pi(J\cup K; x_J, x_K) - \sum_{k\in K} c_k \right]. \qquad (12)
\end{aligned}$$

The above objects are well defined because a straightforward proof by induction implies that the search procedure $((\tilde{E}^t)_{t=1}^{\infty}, (\tilde{q}_i)_{i\in I})$ exists and the function π is well defined.

At each period, given the state variable (J, x_J) for distorted symmetric-information search problem, the seller's alternatives, described relative to any incumbent $i \in J$, are: (i) sell the good to i right now, getting $\mathbb{E}_{x_{-J}} V_i(x)$, (ii) sell to another incumbent $j \neq i$ right now, getting $\mathbb{E}_{x_{-J}} V_j(x)$; (iii) continue search thus getting $\pi_+(J, x_J)$; or (iv) stop and consume the good, getting x_0. The next lemma says that alternative (i) is more likely to be the best option if i's type x_i is high.

Lemma 4. *If $i \in J$, then $\mathbb{E}_{x_{-J}} V_i(x_J, x_{-J})$ and $\pi_+(J, x_J)$ are absolutely continuous functions of x_i; whenever their derivatives exist,*

$$\frac{\partial}{\partial x_i} \mathbb{E}_{x_{-J}} V_i(x_J, x_{-J}) > \frac{\partial}{\partial x_i} \max_{j\in J\setminus i} \mathbb{E}_{x_{-J}} V_j(x_J, x_{-J}) \text{ and} \qquad (13)$$

$$\frac{\partial}{\partial x_i} \mathbb{E}_{x_{-J}} V_i(x_J, x_{-J}) \geq \frac{\partial}{\partial x_i} \pi_+(J, x_J), \qquad (14)$$

and if (14) holds with equality then $\mathbb{E}_{x_{-J}} V_i(x_J, x_{-J}) \geq \pi_+(J, x_J)$.

Proof: By Eq. (2), $\frac{\partial}{\partial x_i} \mathbb{E}_{x_{-J}} V_i(x_J, x_{-J}) = \frac{d}{dx_i} \left[x_i - \frac{1-F_i(x_i)}{f_i(x_i)} \right]$, and, for all $j \neq i$, x_i enters $V_j(x)$ only through $e_{ji}(x_i)$. Hence Assumption 1 implies (13).

To prove (14), we use a revealed-preference argument. Let (J, x_J) be given. Denote

$$\begin{aligned}
\tilde{Q}_k(x_J) &:= \text{the discounted expected probability that } ((\tilde{E}^t)_{t=1}^{\infty}, (\tilde{q}_i)_{i\in I}) \text{ awards} \qquad (15) \\
&\quad\ \text{the good to bidder } k \text{ (who need not has entered) given } (J, x_J).
\end{aligned}$$

Assume that the seller follows the search procedure $((\tilde{E}^t)_{t=1}^{\infty}, (\tilde{q}_i)_{i\in I})$ as if the type of an incumbent $i \in J$ is $\hat{x}_i \neq x_i$. (The seller knows the realized value of x_i because this is a symmetric-information search problem.) Let $\hat{\pi}_+(J, x_J, \hat{x}_i)$ be the expected payoff of the seller from this deviant plan, discounted back to the current period t; x_i enters $\hat{\pi}_+(J, x_J, \hat{x}_i)$ only in the term

$$\mathbb{E}_{x_{-J}} \sum_{k\in I} \delta^{\tau(x_{J\setminus i}, \hat{x}_i, x_{-J})-t} V_k(x_{J\setminus i}, x_i, x_{-J}) \tilde{q}_k(x_{J\setminus i}, \hat{x}_i, x_{-J}).$$

Hence (remember that x_i is one component of x_J)

$$\frac{\partial}{\partial x_i}\hat{\pi}_+(J, x_J, \hat{x}_i) = \tilde{Q}_i(\hat{x}_i, x_{J\setminus i})\frac{d}{dx_i}\left[x_i - \frac{1-F_i(x_i)}{f_i(x_i)}\right] + \sum_{k\in I\setminus i}\tilde{Q}_k(\hat{x}_i, x_{J\setminus i})e'_{ki}(x_i).$$

As $((\tilde{E}^t)_{t=1}^{\infty}, (\tilde{q}_i)_{i\in I})$ solves the dynamic programming problem given the state variable (J, x_J),

$$\pi_+(J, x_J) = \hat{\pi}_+(J, x_J, \hat{x}_i) = \max_{\hat{x}_i} \hat{\pi}_+(J, x_J, \hat{x}_i).$$

Thus, the Milgrom-Segal envelope theorem ([10]) implies that $\pi_+(J, x_J)$ is an absolutely continuous function of x_i and that, whenever its derivative exists,

$$\frac{\partial}{\partial x_i}\pi_+(J, x_J) = \tilde{Q}_i(x_i, x_{J\setminus i})\frac{d}{dx_i}\left[x_i - \frac{1-F_i(x_i)}{f_i(x_i)}\right] + \sum_{k\in I\setminus i}\tilde{Q}_k(x_i, x_{J\setminus i})e'_{ki}(x_i). \tag{16}$$

Thus, Assumption 1 implies (14). If (14) holds with equality, bidder i wins almost surely in subsequent periods if the search were to continue. But since search is costly, this implies in turn that awarding the good to bidder i immediately dominates the option of continuing the search. Hence $\mathbb{E}_{x_{-J}}V_i(x_J, x_{-J}) \geq \pi_+(J, x_J)$. ■

We are now ready to prove the main lemma of this subsection.

Lemma 5. *If Assumption 1 holds, then the search procedure $((\tilde{E}^t)_{t=1}^{\infty}, (\tilde{q}_i)_{i\in I})$ operated under the full disclosure policy satisfies the monotonicity condition (6), i.e., $Q_i(\cdot \mid \mathbb{H}_i)$ is nondecreasing for any information set $\mathbb{H}_i$ in the range of the full-disclosure rule.*

Proof: Let $J \subseteq I$ be the set of incumbents and let $i \in J$. It suffices to prove that $\tilde{Q}_i(x_J)$, defined by (15), is weakly increasing in x_i. We shall prove this claim by induction on the size of $I \setminus J$. The case of $J = I$ follows directly from (13). Pick any $n = 1, 2, \ldots$ and suppose the claim is true if the size of $I \setminus J$ is less than or equal to $n-1$. We shall prove the claim when $I \setminus J$ is of size n. Since the search procedure $((\tilde{E}^t)_{t=1}^{\infty}, (\tilde{q}_i)_{i\in I})$ solves the problem

$$\max\left\{x_0, \mathbb{E}_{x_{-J}}V_i(x_J, x_{-J}), \max_{j\in J\setminus i}\mathbb{E}_{x_{-J}}V_j(x_J, x_{-J}), \pi_+(J, x_J)\right\},$$

Lemma 4 implies that the probability that bidder i wins in the current period is weakly increasing in x_i. The induction hypothesis implies that the probability (discounted back to next period) that he wins later, conditional on him not winning in the current period, is weakly increasing in x_i. Thus, the total discounted winning probability of bidder i is weakly increasing in x_i, as desired. ■

3.1.3 Step three: The irrelevance of disclosure policies

To complete this final step of the proof, we show that the full-disclosure mechanism identified in 3.1.2 is optimal among all incentive feasible mechanisms under any disclosure policy. That is true because the full-disclosure mechanism, by definition of its search procedure, achieves the upper bound identified in 3.1.1 if the mechanism is incentive feasible, and we have proved in 3.1.2 that it is indeed incentive feasible. Thus, disclosure policies are irrelevant because, as we will show below, other disclosure policies cannot yield smaller expected profit than full disclosure.

Formally, denote:

$$\begin{aligned} \Pi_{FD},\ \Pi_{ANY} &:= \text{the seller's optimal discounted expected profit among all incentive} \\ & \text{feasible search mechanisms using, respectively, } FD, ANY; \\ FD, ANY &:= \text{full disclosure, any disclosure policy.} \end{aligned}$$

Since bidders are assumed to be risk neutral, incentive feasibility with a fine disclosure policy implies incentive feasibility with a coarse disclosure policy. (Incentive feasibility means a set of inequalities, one for each event that the disclosure policy can possibly reveal to a bidder; integration across these events preserves the direction of the inequalities.) Hence

$$\Pi_{ANY} \geq \Pi_{FD}.$$

Lemma 3 implies that Π_{ANY} is bounded from above by the optimal discounted expected payoff in the distorted symmetric-information search problem. Since $((\tilde{E}^t)_{t=1}^{\infty}, (\tilde{q}_i)_{i\in I})$ is a solution for this distorted problem, we have

$$\Pi\left((\tilde{E}^t)_{t=1}^{\infty}, (\tilde{q}_i)_{i\in I} \mid (V_i)_{i\in I}\right) \geq \Pi_{ANY}.$$

Lemma 5, coupled with Lemma 2, implies that the search procedure $((\tilde{E}^t)_{t=1}^{\infty}, (\tilde{q}_i)_{i\in I})$ is incentive feasible when it is supplemented by the full disclosure policy. Hence

$$\Pi_{FD} \geq \Pi\left((\tilde{E}^t)_{t=1}^{\infty}, (\tilde{q}_i)_{i\in I} \mid (V_i)_{i\in I}\right).$$

Thus, we obtain the equation that immediately implies the theorem:

$$\Pi_{ANY} = \Pi_{FD} = \Pi\left((\tilde{E}^t)_{t=1}^{\infty}, (\tilde{q}_i)_{i\in I} \mid (V_i)_{i\in I}\right).$$

3.2 An optimal mechanism with private values and no discounting

In this subsection we illustrate Theorem 1 by fully characterizing the optimal mechanism for the case where the bidders have private values ($e_{ij} = 0$ for all $i, j \in I$) and there is no discounting ($\delta = 1$). Since the distributions of bidders' types and the cost of contacting each bidder are not necessarily the same across bidders, the results in this subsection generalize the results in McAfee and McMillan [9] where the bidders are assumed to be ex ante identical.

By Theorem 1, the seller needs to solve the distorted symmetric-information search problem, where his reward from selling the good to a bidder is the bidder's virtual utility. With private values, a bidder's virtual utility depends only on his own type (hence we write $V_i(x_i)$ instead of $V_i(x)$):

$$V_i(x_i) = x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}. \tag{17}$$

This symmetric-information search problem is similar to Weitzman's [16] Pandora problem. In that problem, Pandora searches for the highest reward from n boxes under the assumption that only one box can be opened in any single period and opening each box is costly. Weitzman proved that the solution to Pandora's problem is as follows. First, Pandora computes a cutoff value for each box—this cutoff value depends only on the characteristics of the box itself. Then, if the highest cutoff value falls short of Pandora's initial fallback reward, Pandora does not search and simply gets the fallback reward. Otherwise, Pandora opens the box with the highest cutoff value. In every period, the search continues if the highest cutoff among all closed boxes exceeds the updated fallback reward which is the maximum between the initial fallback reward and the highest reward among all opened boxes. If seach ends, Pandora gets the updated fallback reward.

In our case, the boxes are the potential bidders, and since they are privately informed about their types after being contacted by the seller, the seller's reward from each bidder is the bidder's virtual utility. The initial fallback reward is just x_0. Since we assume away discounting, there is no loss of generality in inviting only one entrant (i.e., "opening one box") in each period as Weitzman's Pandora problem. The relevant cutoff level v_i^* for bidder i is the solution of

$$\mathbb{E}_{x_i}\left[(V_i(x_i) - v_i^*)^+\right] = c_i. \tag{18}$$

This equation is analogous to Eq. (7) in Weitzman [16]. It says that, if the seller's updated fallback payoff is v_i^*, then the seller is indifferent between (i) stopping the search and getting v_i^* immediately, and (ii) contacting bidder i at a cost of c_i and then stopping the search and getting a payoff equal to either $V_i(x_i)$ or v_i^*, whichever higher. When the seller invites bidder i to participate, the bidder makes a report $\hat{x}_i$ and commits to paying, in expected value, an amount $P_i(\hat{x}_i \mid \mathbb{H}_i)$ specified by Eq. (9). The seller keeps searching as long as the highest cutoffs among all bidders who have not yet been contacted exceeds the updated fallback reward. Otherwise the seller stops searching and allocates the good to the bidder with the highest virtual utility among all incumbent bidders, provided that this virtual utility exceeds the seller's value, x_0, or keeps the good otherwise.

Note that, in the above procedure, a participating bidder i can end the search and buy the good immediately if he reports his type as $\hat{x}_i$ such that $V_i(\hat{x}_i)$ is greater than both the virtual utilities of the incumbents and the cutoffs of those not yet contacted. Denote x_i^* for the infimum of such types $\hat{x}_i$. We may call the expected payment $P_i(x_i^* \mid \mathbb{H}_i)$ specified by Eq. (9) the *buy-now price*, as if the bidder were given a buy-now offer to buy the good immediately at the price equal to $P_i(x_i^* \mid \mathbb{H}_i)$.

One way to construct such buy-now prices is as follows: Every participating bidder makes a once-and-for-all bid. Bids are compared to one another in terms of virtual utilities. If a bidder rejects the buy-now offer but eventually wins the good, his payment is equal to the amount which, in terms of virtual utilities, matches the highest losing bid or the seller's own value, whichever is larger. To compute the buy-now prices, suppose the bids collected up to the end of period $t-1$ are $x_1, \ldots, x_{t-1}$, with x_s submitted by the period-s entrant, bidder s. The seller's fallback reward is

$$v^{t-1} := \max\left\{x_0, \max_{s<t} V_s(x_s)\right\}.$$

Suppose that bidder t enters in period t. There are only two possible cases:

i. $v_{t+1}^* \leq v^{t-1}$: the cutoff v_{t+1}^* for contacting the next bidder is below the seller's current fallback reward v^{t-1}, so the search will end no matter what bidder t reports. If bidder t wants to win right now, all that he needs is to submit a bid x_t such that $V_t(x_t) \geq v^{t-1}$, i.e., $x_t \geq V_t^{-1}(v^{t-1})$.[9] Hence the buy-now price for bidder t is $V_t^{-1}(v^{t-1})$.

ii. $v_{t+1}^* > v^{t-1}$: the cutoff for continuing the search is above the seller's current fallback reward, so the search may end or continue, depending on bidder t's report. If his virtual utility is above v_{t+1}^*, then bidder t beats all incumbents and ends the search; yet if his virtual utility is merely equal to v^{t-1}, the search will continue. Thus, the buy-now price, in terms of virtual utilities, should lie between v_{t+1}^* and v^{t-1}. To compute it, first calculate the probability $Q_t^*(z)$ with which bidder t eventually wins in the above search procedure if he reports z, then use

[9] The inverse function V_t^{-1} exists by Assumption 1.

Eq. (9) to obtain the buy-now price as

$$V_t^{-1}(v_{t+1}^*) - \int_{V_t^{-1}(v^{t-1})}^{V_t^{-1}(v_{t+1}^*)} Q_t^*(z)dz.$$

When all bidders have the same distribution for their types (although the search costs of contacting them may be different), it is possible to implement the optimal search mechanism with a sequence of second-price auctions with period-specific reserve prices. The next subsection illustrates this idea by a two-bidders example, whose extension to the n-bidder case is straightforward.

A further special case in which the search costs are also the same across bidders (i.e., the bidders are ex ante identical) has been analyzed by McAfee and McMillan [9]. Since the bidders are ex-ante identical, the cutoffs v_i^* defined by Eq. (18) are identical across bidders. Thus, unlike in most of our paper, the sequence in which the bidders are contacted is irrelevant. Search continues if and only if the highest virtual utility among all incumbent bidders is less than the cutoff. When all bidders have been invited, the good is sold via a conventional second-price auction with reserve.

3.3 A two-bidder example with private values and no discounting

In this subsection we examine a specific two-bidder example that illustrates how the optimal mechanism can be computed. Suppose there are only two potential bidders and $x_i \sim [0,1]$ for $i = 1, 2$. The cost of contacting bidder i is c_i, with $c_1 < c_2 < \frac{1}{4}$. As in Section 3.2, we assume that $e_{12} = e_{21} = 0$ (private values), and $\delta = 1$. Hence, the profit optimal search auction is a Weitzman search procedure with the bidders' values replaced by their virtual utilities

$$V_i(x) = 2x_i - 1. \tag{19}$$

The seller will not sell the good to bidder i if $V_i(x) \leq 0$; hence a bidder with $x_i \leq 1/2$ has a zero probability of winning the good.

Using (18), the cutoffs, v_1^* and v_2^* are defined by

$$\int_{\frac{1+v_i^*}{2}}^{1} [2x - 1 - v_i^*]\, dx = c_i, \qquad \Longleftrightarrow \qquad v_i^* = 1 - 2\sqrt{c_i}.$$

Since $c_1 < c_2 < 1/4$, $v_1^* > v_2^* > 0$, it is optimal to invite bidder 1 in period 1, and if the search continues, to invite bidder 2 in period 2.

Bidder 1 wins the good immediately if $V_1(x_1) \geq v_2^*$, i.e., if $x_1 \geq 1 - \sqrt{c_2}$. Otherwise, the mechanism continues to period 2, and bidder 1 wins if and only if $x_1 > \max\{x_2, 1/2\}$. Since

$x_2 \sim [0,1]$, the probability that bidder 1 wins is 0 if $x_1 < 1/2$, x_1 if $\frac{1}{2} \le x_1 < 1-\sqrt{c_2}$, and 1 if $x_1 \ge 1-\sqrt{c_2}$. Using (9), the expected payment of bidder 1 is

$$p_1(x_1) = \begin{cases} 0, & \text{if} \quad x_1 < \frac{1}{2}; \\ x_1^2 - \int_{\frac{1}{2}}^{x_1} z dz = \frac{x_1^2}{2} + \frac{1}{8}, & \text{if} \quad \frac{1}{2} \le x_1 < 1-\sqrt{c_2}; \\ x_1 - \int_{\frac{1}{2}}^{1-\sqrt{c_2}} z dz - \int_{1-\sqrt{c_2}}^{x_1} dz = \frac{1-c_1}{2} + \frac{1}{8}, & \text{if} \quad x_1 \ge 1-\sqrt{c_2}. \end{cases} \tag{20}$$

Conditional on being invited to participate, bidder 2 wins the good if and only if $V_2(x_2) \ge \max\{V_1(x_1), 0\}$, i.e., $x_2 > \max\left\{x_1, \frac{1}{2}\right\}$. By (9), the expected payment of bidder 2 is

$$p_2(x_2, x_1) = \begin{cases} 0, & \text{if} \quad x_2 \le \max\left\{x_1, \frac{1}{2}\right\}; \\ x_2 - \int_{\max\left\{x_1, \frac{1}{2}\right\}}^{x_2} dz = \max\left\{x_1, \frac{1}{2}\right\}, & \text{if} \quad x_2 > \max\left\{x_1, \frac{1}{2}\right\}. \end{cases}$$

The optimal mechanism can be implemented with the following procedure. Bidder 1 is offered the good at price $(1-c_2)/2 + 1/8$ in period 1; if he rejects the offer, he and bidder 2 participate in period 2 in a second price auction with a reserve price equal to $1/2$.[10] Since both bidders bid their values in the second price auction, bidder 2 wins if and only if $x_2 > \max\{x_1, 1/2\}$ and his payment if he wins is $\max\{x_1, 1/2\}$. Hence, we only need to verify that bidder 1's optimal strategy in our procedure is aligned with the optimal mechanism.

In our procedure, bidder 1 has to pay at least $1/2$ if he wins. Thus, if $x_1 \le 1/2$, he does not want to win, and it is dominant to bid his true type in both periods. In period 2, bidder 1 wins if and only if $x_1 > \max\left\{x_2, \frac{1}{2}\right\}$, and if he wins, he pays $\max\left\{x_2, \frac{1}{2}\right\}$; his expected payment is

$$\int_0^{\frac{1}{2}} \frac{1}{2}\, dx_2 + \int_{\frac{1}{2}}^{x_1} x_2\, dx_2 = \frac{1}{8} + \frac{x_1^2}{2}.$$

This expression is equal to the second line in Eq. (20).

Finally, to verify that bidder 1 accepts the offer in period 1 if and only if $x_1 \ge 1-\sqrt{c_2}$, we note that bidder 1 has two options: (i) agree to pay $\frac{1-c_2}{2} + \frac{1}{8}$ and obtain the good immediately, or (ii) participate in the second price auction with bidder 2. Bidder 1's payoff from option (i) is

$$x_1 - \left(\frac{1-c_2}{2} + \frac{1}{8}\right).$$

With option (ii), bidder 1 wins with probability x_1 and his expected payment is $\frac{x_1^2}{2} + \frac{1}{8}$; his expected payoff is

$$x_1 x_1 - \left(\frac{x_1^2}{2} + \frac{1}{8}\right) = \frac{x_1^2}{2} - \frac{1}{8}.$$

[10] Note that the procedure is equivalent to a sequence of two second price auctions, with period-specific reserve prices. The first auction in the sequence is degenerate as only bidder 1 participates.

Comparing bidder 1's expected payoff under the two options shows that he will choose option (i) if and only if $x_1 \geq 1 - \sqrt{c_2}$. Hence bidder 1's strategy is consistent with the optimal mechanism.

4 Properties of profit maximizing search mechanisms

In standard auction theory, asymmetric information leads to inefficiencies in the form of no trade in some states of nature, and, sometimes, biased allocations. In our search-theoretic framework, asymmetric information leads to a third form of inefficiency: inefficient search procedures. In Subsection 4.1, we show that the optimal mechanism may completely exclude some bidders who would be invited to participate in the (socially) efficient mechanism. In Subsection 4.2, we show an opposite effect: the optimal mechanism gives the seller an excessive incentive to search relative to the efficient mechanism. In Subsection 4.3, we show that the order in which bidders are approached need not be the same as the efficient mechanism.

In order to explore these kinds of inefficiency, we shall consider the private values case where $e_{ij} = 0$ for all $i, j \in I$ and there is no discounting.[11] Given these assumptions, the seller's optimal search mechanism is fully characterized in Subsection 3.2. In particular, the optimal search procedure is generated by Weitzman's solution with the cutoffs being implicitly defined by Eq. (18). By contrast, the cutoffs for a (socially) efficient procedure are defined by the equation

$$\mathbb{E}_{x_i} \left[x_i - x_i^* \right]^+ = c_i. \tag{21}$$

This difference arises because the payoff in the associated search problem is measured in virtual utilities in the former mechanism and is measured in actual utilities in the latter. In Crémer, Spiegel, and Zheng [6] we proved that the efficient search procedure can always be implemented by a perfect Bayesian equilibrium.[12]

4.1 Fewer participants

Because a bidder's actual value exceeds his virtual utility, the benefit of inviting a bidder to participate is lower in an optimal search mechanism for the seller than it is in an efficient mechanism for the economy if the fallback payoffs are the same. In fact, the optimal mechanism may completely

[11] By continuity, the inefficiencies that we identify still hold if these assumptions are slightly relaxed.

[12] Although the participation constraint is ex ante in that paper, its efficiency result is applicable here because interim participation constraints can always be satisfied by transfers from the seller.

exclude a bidder even before the search begins, even though that bidder has a positive probability of participation in an efficient mechanism. The consequences of this fact are described in the following proposition.

Proposition 1. *¿From the standpoint of period 1, every bidder i's probability of participation in a socially efficient mechanism is positive if his probability of participation in an optimal mechanism is positive, but the converse is not necessarily true.*

Proof: For $z \le \overline{x}_i$, $\mathrm{E}_{x_i}\left[V_i(x_i) - z\right]^+$ and $\mathrm{E}_{x_i}\left[x_i - z\right]^+$ are strictly decreasing functions of z and $\mathrm{E}_{x_i}\left[V_i(x_i) - z\right]^+ \le \mathbb{E}_{x_i}\left[x_i - z\right]^+$. Hence, $v_i^* < x_i^*$ for all $i \in I$. The proof is completed by noting that a bidder i has a positive probability of participating in the socially efficient mechanism if $x_i^* > x_0$ and a positive probability of participation in the optimal mechanism *only* if $v_i^* \ge x_0$. ■

4.2 Longer search

As virtual utilities are below actual utilities, the seller's fallback value in the optimal mechanism is smaller than his fallback value in a socially efficient mechanism. This leads to an effect opposite to the previous one, as the lower fallback value makes it more attractive to continue the search. A simple case for this effect is that bidders' types are *i.i.d.*, so their virtual utility functions are the same, say V, though their participation costs may still be different. While the cost of an additional searching period is the same in both the efficient and optimal mechanisms, the gains are different. To see that, suppose that an additional search increases the highest reported value and hence the social surplus by Δx. The resulting effect on the seller's revenue, measured in virtual utilities, is approximately $V'(x)\Delta x$. Under Assumption 1, $V'(x) \ge 1$. Thus, other things equal, the seller is more willing to continue searching than a social planner would.

Proposition 2. *Assume that types x_i are identically distributed across bidders, with V denoting their common virtual utility function (though their participation costs c_i may be different). If Assumption 1 holds and if $v_i^* > x_0$ for all $i \in I$, then an optimal search lasts at least as long as, and with a positive probability strictly longer, a socially efficient search.*

Proof: First, recall that $\overline{x}_i$ is the upper bound of the support of x_i and let

$$\phi(z) := \int_z^{\overline{x}_i} (x_i - z) dF_i(x_i); \quad \varphi(z) := \int_z^{\overline{x}_i} (V_i(x_i) - V_i(z)) dF_i(x_i).$$

The solution of $\varphi(z) = c_i$ is $V_i^{-1}(v_i^*)$. By assumption $V' > 1$, $\varphi' < \phi' < 0$ throughout their common domain. Then the fact that $\phi(\overline{x}_i) = 0 = \varphi(\overline{x}_i)$ implies $V^{-1}(v_i^*) > x_i^*$ for all $i \in I$.

Second, the order of entry is the same in both mechanisms: with i.i.d. bidders, $x_i^* > x_j^*$ if and only if $c_i < c_j$ if and only if $v_i^* > v_j^*$. Thus, we can relabel the bidders so that $v_1^* \geq v_2^* \geq \cdots \geq v_n^*$ and $x_1^* \geq x_2^* \geq \cdots \geq x_n^*$. Because $v_i^* > x_0$, the optimal mechanism invites bidder 1 to participate. We show that it continues from period t to period $t+1$ with a higher probability than the socially efficient procedure. To see that, let $(x_1, \ldots, x_t)$ be the sequence of realized values up to period t. If the efficient search continues to period $t+1$, then $\max\{x_0, x_1, \ldots, x_t\} < x_{t+1}^*$; because $V^{-1}(v_i^*) > x_i^*$ and $v_i^* > x_0$,

$$v_{t+1}^* > \max\{x_0, V_1(x_1), \ldots, V_t(x_t)\}.$$

Hence the optimal search continues to period $t+1$. Thus, it continues whenever the efficient procedure continues. The converse, however, is false: when

$$x_{t+1}^* < \max\{x_1, \ldots, x_t\} < V^{-1}(v_{t+1}^*),$$

which occurs with a positive probability, the efficient search stops while the optimal search procedure continues. This proves the proposition. ■

4.3 Inefficient order of entry

Determined by different sets of cutoffs, the order in which the bidders enter the optimal mechanism may differ from the order in the socially efficient mechanism. The following example shows this distortion with two bidders. The seller's value, x_0, is zero; bidder 1's type x_1 is uniformly distributed on $[\underline{x}_1, \overline{x}_1]$ and bidder 2's type x_2 is drawn from an exponential distribution $F_2(x_2) := 1 - \exp(-\lambda x_2)$. Given these assumptions, the virtual utility functions and cutoffs are

$$\begin{aligned} V_1(x_1) &= 2x_1 - \overline{x}_1; \qquad & V_2(x_2) &= x_2 - 1/\lambda; \\ x_1^* &= \overline{x}_1 - \sqrt{2c_1(\overline{x}_1 - \underline{x}_1)}; \qquad & x_2^* &= -\ln(\lambda c_2)/\lambda; \\ v_1^* &= \overline{x}_1 - 2\sqrt{c_1(\overline{x}_1 - \underline{x}_1)}; \qquad & v_2^* &= -\left(1 + \ln(\lambda c_2)\right)/\lambda. \end{aligned}$$

Since $x_i^* > v_i^*$, there exist two numbers a and b such that $v_1^* < a < b < x_1^*$. Let $\lambda := 1/(b-a)$ and $c_2 := \exp(-\lambda b)/\lambda$. Then $x_1^* > x_2^*$ and $v_1^* < v_2^*$: bidder 2 enters first in the optimal mechanism, whereas bidder 1 enters first in the socially efficient mechanism.

4.4 Delayed participation of influential bidders

When bidders' values are interdependent, we can address the following question: If bidder i's type has a stronger influence on the valuations of other bidders than bidder j's type, should the seller let i enter before j or vice versa? For simplicity, we address this question under the assumption that, for all $j \in I$, there is a number α_j such that $e_{ij}(x_j) = \alpha_j x_j$ for all x_j and $i \neq j$, with $\alpha_i \in [0, 1)$. We can therefore regard bidders with higher α's as more influential. We show that the more influential bidders will enter later than less influential bidders. As we will see, this is a property linked to the sequential nature of the search mechanism, and not, as in Subsections 4.1 to 4.3, a distortion due to asymmetry of information.

Proposition 3. *Assume $V_i \geq 0$ and $\underline{x}_i \geq 0$ for all i, and $e_{ji}(x_i) = \alpha_i x_i$ for all $j \neq i$. Also assume that search costs and type-distributions are identical across $i \in I$. Then the larger α_i is, the later and less probable is i's entry in an optimal or a socially efficient search mechanism.*

Proof: For every $i \in I$ and every $x_i \in X_i$, let

$$W_i(x_i) := x_i - \frac{1 - F_i(x_i)}{f_i(x_i)} - \alpha_i x_i.$$

Note that $V_i(x) = W_i(x_i) + \sum_{j \in I} \alpha_j x_j$. Thus, for any state variable (J, x_J),

$$\mathbb{E}_{x_{-J}} V_i(x_J, x_{-J}) > \mathbb{E}_{x_{-J}} V_j(x_J, x_{-J}) \Longleftrightarrow W_i(x_i) > W_j(x_j). \tag{22}$$

Since $V_i \geq 0$ by assumption, the optimal mechanism never results in no sale. Hence (22) implies that the search procedure is equivalent to a standard Weitzman seach with payoff from search being $W_i(x_i)$. In this search procedure, the cutoffs w_i^* are implicitly defined by

$$\mathbb{E}_{x_i} \left[W_i(x_i) - w_i^*\right]^+ = c_i.$$

$\mathbb{E}_{x_i} \left[W_i(x_i) - w_i^*\right]^+$ is strictly decreasing in w_i^* and $W_i(x_i)$ is strictly decreasing α_i (since $x_i \geq 0$ by assumption). Thus, w_i^* is strictly decreasing in α_i, as claimed.

The proof for the socially efficient mechanism is analogous and can be deduced by simply setting $W_i(x_i) := x_i - \alpha_i x_i$ for every $i \in I$ and every $x_i \in X_i$. ∎

The basic intuition for Proposition 3 is as follows: if there is a very influential bidder, a change in his type will increase the value of other bidders nearly as much as it increases his own value. Hence, inviting this bidder to participate in the mechanism early on will not reduce the set of states of nature in which it would also be optimal to invite other bidders to participate and will therefore not save on search costs. To illustrate, assume that there are two bidders, with $\alpha_1 = 1$, $\alpha_2 = 0$ and $x_0 = 0$. Bidder 1 is therefore "influencial" as a change in x_1 has a positive effect on bidder 2's value but not vice versa. Since bidder 2's value, $x_2 + x_1$, always exceeds bidder 1's value, x_1, it is clear that if the good is allocated, it is certainly allocated to bidder 2; thus, only bidder 2 will be invited to participate. The influential bidder, bidder 1, is excluded.

5 Conclusion

We have studied a single unit auction environment in which the set of bidders is endogenously determined through a dynamic search process. The distinctive feature of our model is that it allows for discounting and asymmetry and value-interdependency across potential bidders. With asymmetry and interdependency, the sequence in which bidders are contacted is important. With discounting, some potential bidders may need to be contacted at the same time.

We showed that an optimal mechanism for the seller is equivalent to an optimal search with symmetric information, where the utilities are replaced by virtual utilities. That is, the seller conducts a costly search for the bidder with the highest virtual utility. In standard auction theory, the information rents that the seller concedes to the bidders create inefficiencies in the form of no trade in some states of nature and, sometimes, biased allocations. Our search-theoretic framework gives rise to a third form of inefficiency: inefficient search procedures. In the case of private values with no discounting, this inefficiency results in fewer participants, longer search conditional on the same set of participants, and inefficient sequence of entry, relative to the socially efficient mechanism.

Appendix: A formal definition of revelation search mechanisms

In this appendix, we complement the informal definition of a sequential mechanism of 2.3 by a formal definition, couched in decision theoretical language. This enables us to provide a formal link between 2.3 and the definitions introduced in 2.4. To this end we first define a *decision tree* for the seller which is associated with his search problem. Then we define a search procedure as a transition function on the seller's decision tree, and payment scheme and disclosure policy as plans contingent on the nodes.[13]

The initial node of the decision tree, d^0, is a *decision node*. The immediate successors of d^0 are (d^0, E^1), where $E^1 \subseteq I$ is the set of period-1 entrants at that node; if $E^1 = \varnothing$, then the seller keeps the good without conducting any search. Each (d^0, E^1) is a *history node*. The immediate successors of (d^0, E^1) are (d^0, E^1, x^1), where $x^1 \in \times_{i\in E^1} X_i$ is the profile of the realized values of period-1 entrants at that node. Each (d^0, E^1, x^1) is a decision node at period 1.

Let $d^t := (d^0, E^1, x^1, \ldots, E^t, x^t)$ be a period-t decision node. The immediate successors of d^t are either (d^t, w) or (d^t, E^{t+1}), where $w \in \cup_{s=1}^t E^s \cup \{\text{seller}\}$ is the winner of the good (which could be the seller), and $E^{t+1} \subseteq I \setminus \cup_{s=1}^t E^s$, with[14] $E^{t+1} \neq \varnothing$, is the set of period-$(t+1)$ entrants. Each (d^t, w) is a *terminal node* specifying who gets the good, while each (d^t, E^{t+1}) is a *nonterminal history node* specifying the potential bidders who are invited to participate in period $t+1$. Note that if $I \setminus \cup_{s=1}^t E^s = \varnothing$, then the immediate successors of d^t are all terminal. The immediate successors of any nonterminal history node (d^t, E^{t+1}) are (d^t, E^{t+1}, x^{t+1}), where $x^{t+1} \in \times_{i\in E^{t+1}} X_i$ is the profile of the realized values of period-$(t+1)$ entrants at that node, and (d^t, E^{t+1}, x^{t+1}) is a period-$(t+1)$ decision node. Hence a decision tree is recursively defined.

Let D be the set of all decision nodes. For any $d^t \in D$, let

$$(d^t)_+ := \left\{(d^t, w) : w \in \cup_{s=1}^t E^s \cup \{\text{seller}\}\right\} \cup \left\{(d^t, E^{t+1}) : \varnothing \neq E^{t+1} \subseteq I \setminus \cup_{s=1}^t E^s\right\}$$

be the set of immediate successors of d^t, and $(D)_+$ be the set of successors of decision nodes. With ΔS denoting the set of lotteries on the outcome set S, a *search procedure* is a transition function

$$\sigma : D \longrightarrow \Delta(D)_+ \text{such that} \sigma(d^t) \in \Delta(d^t)_+.$$

[13]This type of formalism was first developed by Zheng [19].

[14]The assumption $E^{t+1} \neq \varnothing$ means that the seller does not stop searching for some periods and then start searching again. This is trivially true for any optimal mechanism when $\delta < 1$. When $\delta = 1$, there exists an optimal mechanism that satisfies this assumption.

A *payment scheme* is a function p on the set $(D)_+$ that assigns to each immediate successor of $(d^0, E^1, \ldots, E^t, x^t)$ a profile $(p_i^t)_{i \in \cup_{s=1}^t E^s} \in \mathbb{R}^{\cup_{s=1}^t E^s}$ of payments, where p_i^t is a participating bidder i's payment to the seller delivered in period t (nonparticipants make zero payments).

A *disclosure policy* is a function H ("history") on the set of nonterminal history nodes such that, at each such node (d^{t-1}, E^t), H assigns to each period-t entrant $i \in E^t$ an information set $H_i(d^{t-1}, E^t)$ that contains this node. A period-t entrant i's *knowledge* consists of his realized value, all the rules of the search mechanism, and the fact that he is at a node belonging to $H_i(d^{t-1}, E^t)$. A bidder makes the participation decision only after he has acquired this knowledge.

In a *non-disclosure* policy, for any entrant $i \in E^t$, $H_i(d^{t-1}, E^t)$ consists of all the nonterminal history nodes $(d^0, E^1, \ldots, E^s)$, with $s = 1, 2, \ldots$, such that $i \in E^s$ (so i cannot tell any two of these nodes apart); hence i does not even know which period he is in, let alone the past sequence of entry or the reports from previous entrants. In a *full disclosure* policy, for any entrant $i \in E^t$, $H_i(d^{t-1}, E^t) = \{(d^{t-1}, E^t)\}$; hence i knows the entire past history, including the reports from previous entrants.

A revelation search mechanism is a triplet (σ, p, H), consisting of a search procedure σ, payment scheme p, and disclosure policy H. Note that since σ is a transition function, this definition allows randomization of the sequence of search, payment scheme, and disclosure policy.

From the above definition we can derive the notations used in the main text. Let $\tilde{\omega}$ be a random vector with support Ω such that, conditional on any decision node $d^t \in D$, any realization ω of $\tilde{\omega}$ uniquely determines an immediate successor $s(\omega)$ of d^t, and the probability with which a successor is $s(\omega)$ is equal to the probability governed by the transition function $\sigma(d^t)$. A *realized state* $\tilde{x}$ is a pair of bidder-type profile and realization of $\tilde{\omega}$:

$$\tilde{x} := (x, \omega) := ((x_i)_{i \in I}, \omega) \in X \times \Omega.$$

If every invited bidder participates and is truthful, then any realized state x uniquely determines a complete history of the search, i.e., a terminal node

$$(d^\tau, w) = (d^0, E^1, x^1, \ldots, E^\tau, x^\tau, w),$$

where τ is the terminal period at which the mechanism ends, $(E^1, \ldots, E^\tau)$ is the sequence of entry, and w is the winner of the good (possibly the seller). Abusing notations, denote the terminal period by $\tau(\tilde{x})$, the good's winner $w(\tilde{x})$, and the sequence of entry by $(E^1(\tilde{x}), \ldots, E^{\tau(\tilde{x})}(\tilde{x}))$. If bidder i enters at period t during this search history, i's information set is uniquely determined to be some $H_i(d^{t-1}(\tilde{x}), E^t(\tilde{x}))$; abusing notations, denote this t by $t_i(\tilde{x})$, and this information set by $H_i(\tilde{x})$.

Thus, given any revelation search mechanism (σ, p, H), if every invited bidder participates and is truthful, then any realized state $\tilde{x} \in X \times \Omega$ uniquely determines the following objects:

$$\begin{aligned}
E^t(\tilde{x}) &:= \text{the set of potential bidders who enter the mechanism in period } t;\\
1_i(\tilde{x}) &:= \text{the indicator function for "player } i \text{ finally owns the good"};\\
H_i(\tilde{x}) &:= \text{the information set for bidder } i \text{ when } i \text{ enters};\\
t_i(\tilde{x}) &:= \text{the period at which player } i \text{ enters the mechanism};\\
\tau(\tilde{x}) &:= \text{the period at which the search terminates}.
\end{aligned}$$

All our calculations remain intact if these notations replace their counterparts ($E^t(x)$, $q_i(x)$, etc.) in the main text. To keep the exposition simple, we prefer however to use in the main text the notations introduced in Section 2.4. Moreover we assume in the main text that the search procedure is not randomized; this assumption is harmless however because every player is risk neutral.

References

[1] Dirk Bergemann and Martin Pesendorfer. Information structures in optimal auctions. Cowles Foundation Discussion Paper 1323, August 2001.

[2] Dirk Bergemann and Juuso Välimäki. Information acquisition and efficient mechanism design. *Econometrica*, 70(3):1007–1033, May 2002.

[3] Dirk Bergemann and Juuso Välimäki. Information in mechanism design. Cowles Foundation Discussion Paper 1532, August 2005.

[4] Roberto Burguet. Optimal repeated purchases when sellers are learning about costs. *Journal of Economic Theory*, 68:440–455, 1996.

[5] Olivier Compte and Philippe Jehiel. On the virtues of the ascending price auction: New insights in the private value setting. Mimeo, C.E.R.A.S.-E.N.P.C., C.N.R.S., December 2000.

[6] Jacques Crémer, Yossef Spiegel, and Charles Z. Zheng. Optimal selling mechanisms with costly information acquisition. Working paper, Institut d'Economie Industrielle (IDEI), Toulouse, August 12, 2003.

[7] Shmuel Gal, Michael Landsberger, and Arkadi Nemirovski. Participation in auctions. Mimeo, University of Haifa, Haifa, Israel, 2003.

[8] Dan Levin and James L. Smith. Equilibrium in auctions with entry. *American Economic Review*, 84(3):585–599, June 1994.

[9] R. Preston McAfee and John McMillan. Search mechanisms. *Journal of Economic Theory*, 44:99–123, 1988.

[10] Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, March 2002.

[11] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, February 1981.

[12] Wolfgang Pesendorfer and Asher Wolinsky. Second opinion and price competition: Inefficiency in the market for expert advice. *Review of Economic Studies*, 70(243):417–437, April 2003.

[13] Leonardo Rezende. Mid-auction information acquisition. Mimeo, Department of Economics, Stanford University, October 2000.

[14] Mark Stegeman. Participation costs and efficient auctions. *Journal of Economic Theory*, 71:228–259, 1996.

[15] Tara Vishwanath. Parallel search for the best alternative. *Economic Theory*, 2:495–507, 1992.

[16] Martin L. Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3):641–654, May 1979.

[17] Asher Wolinsky. Competition in a market for informed experts' services. *RAND Journal of Economics*, 24(3):380–398, Autumn 1993.

[18] Lixin Ye. Optimal auctions with endogenous entry. *Contributions to Theoretical Economics*, 4(1):Article 8, 2004.

[19] Charles Z. Zheng. VCG mechanisms in dynamic processes. Mimeo, http://faculty-web.at.northwestern.edu/economics/zheng/research/current/vcg/, Oct. 2005.

Index