Available online at www.sciencedirect.com
ScienceDirect
Journal of Economic Theory 134 (2007) 226–248
JOURNAL OF Economic Theory
www.elsevier.com/locate/jet

Optimal search auctions ☆

Jacques Crémer[a], Yossi Spiegel[b], Charles Zhoucheng Zheng[c,*]

[a] *IDEI-GREMAQ, Université de Toulouse, France*
[b] *Recanati Graduate School of Business Administration, Tel Aviv University, Ramat Aviv, Tel Aviv 69978, Israel*
[c] *Department of Economics, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208, USA*

Received 17 June 2004; final version received 3 March 2006
Available online 3 May 2006

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.

JEL classification: D44; D82; D83

Keywords: Optimal auction; Participation cost; Search procedure

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 such 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.

☆ This version supersedes a draft titled "Optimal Auctions from a Search-Theoretic Viewpoint".

* Corresponding author.
E-mail addresses: jacques@cremeronline.com (J. Crémer), spiegel@post.tau.ac.il (Y. Spiegel), czheng@northwestern.edu (C.Z. Zheng)
URLs: http://www.tau.ac.il/~spiegel (Y. Spiegel), http://faculty-web.at.northwestern.edu/economics/zheng/ (C.Z. Zheng).

0022-0531/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.jet.2006.03.003

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 [12]. 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 prospective 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 prospective 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 incurring 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.

In Section 3.2, this theorem is applied to the private-value case without discounting. The optimal mechanism there can be interpreted as a sequence of Myerson's [12] optimal auctions, with one additional bidder in each period. If all potential bidders have the same distribution of their values (though the costs of contacting them may be different), the optimal mechanism is simply a sequence of second-price auctions with reserve prices that decline over time.

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 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.[1] Myerson [12] 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 [10] 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

[1] 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 [17] and Vishwanath [16] considers 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.

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 [10] 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 [9], and Ye [19]. (See Bergemann and Välimäki [3] 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 [14], who analyze the effect of information acquisition conducted during ascending price auctions.[2]

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.[3] Second, although we consider an auction environment, our framework can be easily modified to a procurement environment in which a procurer wishes to

[2] There are a few other interesting papers on specific auction formats with information acquisition or costly participation, including Gal, Landsberger and Nemirovski [7], Stegeman [15], Pesendorfer and Wolinsky [13], and Wolinsky [18]. However, unlike in our paper, the participation or information acquisition decisions in these papers are not coordinated by a principal.

[3] 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.

procure an indivisible good from a set of potential suppliers. If the procurer's 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.[4]

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 [12], 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 $\left\{e_{ij}(\cdot)\right\}_{j \in I \setminus i}$. 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 $t' \geqslant s$, 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 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.[5] 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

[4] 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.

[5] 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.

are invited. Depending on the information disclosure policy that the seller adopts as part of the mechanism, each entrant in periods 1 and 2 is told none, or part of, or all of the messages sent by period-1 entrants. Therefore, a period-2 entrant need not even know how many periods the search has lasted.[6] Given the disclosed information, each period-2 entrant decides whether to participate; if he does, he sends a message. If the rules of the mechanism allow it, a period-1 entrant can submit a new 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.

2.4. Revelation search mechanisms

A *revelation search mechanism* is a search mechanism in which (i) each bidder i's message space is i's type space, (ii) each entrant is disclosed information only when he is being contacted, and (iii) each entrant sends a message once and only once, immediately after his information disclosure stage. 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 standard one and hence is omitted. In a general search mechanism, a participating bidder may submit messages in several periods and hence his equilibrium strategy is a mapping from his type to a complete contingency plan. However, there is no loss of generality in suppressing the contingency plan and replacing it with a one-shot message. The principal, with full commitment power, can tell the bidder: "Tell me your type and I will play for you in the multistage mechanism according to the contingency plan that you would have chosen yourself."

Lemma 1 (*Revelation principle for search mechanisms*). *For any equilibrium feasible search mechanism, there exists an incentive feasible revelation search mechanism that replicates its equilibrium outcome.*

Given a revelation search mechanism, suppose that the profile of realized types is $x \in X$ and every invited bidder participates and is truthful. Then the mechanism induces the following objects (the formal definition is in the Appendix):

$E^t(x) :=$ the set of potential bidders who enter the mechanism in period t.
$q_i(x) :=$ the probability with which player i (bidder or seller) consumes the good.
$p_i(x) :=$ the total payment made by bidder i discounted back to the period at which i enters the mechanism.
$H_i(x) :=$ the smallest subset of $\times_{j\in I\setminus i} X_j$ which bidder i knows, when i enters the mechanism, that contains the profile of all other players' realized types.
$t_i(x) :=$ the period at which player i enters the mechanism ($i \in E^{t_i(x)}(x)$).
$\tau(x) :=$ the period at which the search ends.
$I^t(x) := \bigcup_{s=1}^{t} E^s(x) =$ the set of incumbents at the end of period t.

[6] By contrast, a period-1 entrant knows that he is in the first cohort. That is because the set of period-1 entrants is part of the rules of the mechanism which are explained to each bidder when he is contacted. This feature, however, is not essential: as explained in the Appendix, the seller can randomize the search procedure, in which case even a period-1 entrant need not know he is in the first cohort.

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 $t \geqslant 1$, 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 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]
(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 *full disclosure* policy, for any profile x of realized types and in any period t, every entrant knows the entire sequence $(E^s(x))_{s=1}^t$ of entry up to that point, as well as the reported type x_j of every incumbent $j \in I^{t-1}(x)$. In a *non-disclosure* policy, an entrant only learns that he was invited to participate but he does not learn the history of the mechanism up to that point, nor even how many periods the mechanism was played.

Whatever the disclosure policy is, the bidders are told about 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 and types are independent. 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

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

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

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

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. A search mechanism is *optimal* (profit maximizing) if the seller's discounted expected profit is maximized over all search mechanisms subject to interim participation and incentive compatibility constraints. A search procedure is *optimal* if it is the search procedure of an optimal search mechanism.

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 $\frac{d}{dx_i}\left(x_i - \frac{1-F_i(x_i)}{f_i(x_i)}\right) > e'_{ji}(x_i) \geqslant 0$ 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. No matter how complicated a search procedure is, what matters from each bidder's viewpoint 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 (recall that each bidder 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 type 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 seller's surplus extracts from each bidder i is uniquely determined by the bidder's type x_i and is independent of the disclosure policy.[9]

3.1. *The proof of the theorem*

We prove Theorem 1 in three steps. First, in Section 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 Section 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 Section 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: 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

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

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

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

[9] Gershkov and Szentes [8] study a dynamic mechanism design problem in which committee members acquire costly information in order to make a collective decision, and show that full disclosure is not optimal. Their problem however differs from ours in many important dimensions (e.g., there are no monetary transfers and the cost of information acquisition is borne by the agents). Currently, there still does not exist a comprehensive framework that unifies the various strands of the literature on dynamic mechanism design with costly information acquisition that makes it possible to identify the general circumstances under which full disclosure is or is not optimal.

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 [12, 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 *(i.e.,* $\{H_i(x) : x \in X\}$*):*

$$\text{the function } Q_i(\cdot \mid \mathbb{H}_i) \text{ is nondecreasing,} \tag{6}$$

$$u_i(x_i \mid x_i, \mathbb{H}_i) = u_i(\underline{x}_i \mid \underline{x}_i, \mathbb{H}_i) + \int_{\underline{x}_i}^{x_i} Q_i(z \mid \mathbb{H}_i)\, dz, \tag{7}$$

$$u_i(\underline{x}_i \mid \underline{x}_i, \mathbb{H}_i) = 0. \tag{8}$$

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 equal to*

$$\Pi((E^t)_{t=1}^{\infty}, (q_i)_{i\in I} \mid (V_i)_{i\in I}),$$

defined by (1) *with* u_i *replaced by* V_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 Eq. 9, 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 [12, Lemma 3]), this is equal to

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

where we used Eqs. (2)–(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 types: $t = t_i(x)$.

Thus, viewed from period 1, the seller's expected profit extracted from bidder i is

$$\mathbb{E}_t\left[\mathbb{E}_{\mathbb{H}_i}\left[\mathbb{E}_x\left[\delta^{\tau(x)-1}q_i(x)\left(V_i(x)-x_0\right)-\delta^{t-1}c_i\right]\middle|\,\mathbb{H}_i=H_i(x)\right]\middle|\,t=t_i(x)\right]$$

$$=\mathbb{E}_x\left[\delta^{\tau(x)-1}q_i(x)\left(V_i(x)-x_0\right)-\sum_{t:\,t=t_i(x)}\delta^{t-1}c_i\right]. \tag{10}$$

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

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

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

$J :=$ the set of incumbents, bidders who have entered,

$x_J := (x_i)_{i\in J}, \quad x_{-J} := (x_i)_{i\notin J},$

$\pi(J, x_J) :=$ the seller's optimal discounted expected payoff given the state variable (J, x_J),

$\pi_+(J, x_J) :=$ the seller's optimal discounted expected payoff from continuing the search given (J, x_J)

$$= \delta \max_{K\subseteq I\setminus J}\left[\mathbb{E}_{x_K}\pi(J\cup K; x_J, x_K) - \sum_{k\in K}c_k\right]. \tag{12}$$

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,*

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

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

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

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

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

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),

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

Thus, the Milgrom–Segal envelope theorem [11] 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 $\mathbb{E}_{x_{-J}} V_i(x_J, x_{-J}) \geqslant \pi_+(J, x_J)$. $\square$

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 claimed. □

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 Section 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 Section 3.1.1 if the mechanism is incentive feasible, and we have proved in Section 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

$\Pi_{\mathrm{FD}}, \Pi_{\mathrm{ANY}} :=$ the seller's optimal discounted expected profit among all incentive feasible search mechanisms using, respectively, "full disclosure" (FD), "any disclosure policy" (ANY).

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

$$\Pi_{\mathrm{ANY}} \geqslant \Pi_{\mathrm{FD}}.$$

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

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

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

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

Thus, we obtain the equation that immediately implies the theorem

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

3.2. An optimal mechanism with private values and no discounting

In this subsection we illustrate Theorem 1 for the case where bidders have private values ($e_{ij} = 0$ for all $i, j \in I$) and there is no discounting ($\delta = 1$). We shall show that the optimal search procedure can be implemented by a sequence of Myerson's [12] optimal auctions, with an additional bidder joining the competition in each period if the good is not yet sold. In the special case where bidders have the same distribution of values, this sequence of auctions is simply a sequence of second-price auctions with period-specific reserve prices that decline over time.

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 [10] 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 [17] 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 ex ante 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 search ends, Pandora gets the updated fallback reward.

In our distorted symmetric-information search problem, the boxes are the potential bidders, and 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

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

This equation is analogous to Eq. (7) in Weitzman [17]. 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 is higher.

With the cutoffs computed, the optimal search mechanism is: relabel the potential bidders so that $v_1^* \geqslant \cdots \geqslant v_n^* \geqslant v_{n+1}^* := -\infty$, n being the total number of potential bidders. If $0 \geqslant v_1^*$, quit. Else invite bidder 1. Every invited bidder is asked to make 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). If $(\hat{x}_s)_{s=1}^{t-1}$ are the invited bidders' reports up to the end of period $t-1$, let

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

If $v^{t-1} < v_t^*$, continue to period t and invite bidder t. If $v^{t-1} \geqslant v_t^*$, the seller stops searching and sells the good to a bidder $s \in \{1, \ldots, t-1\}$ for whom $V_s(\hat{x}_s) = v^{t-1}$, if such bidder exists, or keeps the good otherwise.

Note that bidder t can end the search and buy the good immediately if he reports his type as $\hat{x}_t$ such that $V_t(\hat{x}_t) \geqslant \max\{v^{t-1}, v_{t+1}^*\}$. In that case, Eq. (9) implies that the payment made by bidder t is equal to[10]

$$\begin{aligned} V_t^{-1}(\max\{v^{t-1}, v_{t+1}^*\}) &= \max\left\{V_t^{-1}\left(\max_{s<t} V_s(\hat{x}_s)\right), V_t^{-1}(\max\{x_0, v_{t+1}^*\})\right\} \\ &= \max\left\{V_t^{-1}\left(\max_{s<t} V_s(\hat{x}_s)\right), r_t^t\right\}, \end{aligned} \tag{19}$$

where

$$r_t^t := V_t^{-1}(\max\{x_0, v_{t+1}^*\}). \tag{20}$$

Thus, bidder t is essentially competing in Myerson's [12] optimal auction:

(a) he wins immediately if and only if (i) the virtual utility $V_t(\hat{x}_t)$ of his bid is the highest among all invited bidders $1, \ldots, t$ and (ii) $\hat{x}_t$ is at least as high as the reserve price r_t^t;
(b) if he wins immediately, his payment is equal to either (i) the minimum bid that defeats all other invited bidders in virtual utilities or (ii) the reserve price r_t^t, whichever is higher;
(c) if he does not win immediately, bidder t may still win at the end of the search; at that point, his winning event is (i) $V_t(\hat{x}_t)$ is highest among all the bidders that have been invited and (ii) $\hat{x}_t$ matches the reserve price $V_t^{-1}(x_0)$ set at the final period of the search.[11]

Note that bidder t's incentive is unchanged if the bidders invited in previous periods are allowed to resubmit their bids so that $(\hat{x}_s)_{s=1}^{t-1}$ records their most recent instead of old bids. These earlier bidders are truth-telling in the resubmission if any inconsistency between their new and old bids is forgiven and if their new bids are subject to the optimal search procedure, so that the winning event for bidder s is (i) the virtual utility of her new bid is greater than the virtual utilities of the most recent bids from all other invited bidders, $i = \{1, \ldots, t\} \setminus \{s\}$, and (ii) the new bid is at least as high as a bidder-specific reserve price r_s^t:

$$r_s^t := V_s^{-1}(\max\{x_0, v_{t+1}^*\}), \tag{21}$$

which weakly decreases as t increases, as v_{t+1}^* weakly decreases in t.

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

[11] If the search stops at the end of period T, then $v^T \geqslant v_{T+1}^*$ (including the case $T = n$ since $v_{n+1}^* = -\infty$), hence bidder t's winning event in the final period is $V_t(\hat{x}_t) \geqslant v^T$.

Hence the optimal mechanism has the following interpretation: if the search continues to period t, invite bidder t and hold a Myerson auction with period- and bidder-specific reserve prices r_s^t among all the invited bidders $s \in \{1, \ldots, t\}$.

In the special case where all bidders have the same distribution for their types (though the search costs of contacting them may be different), the virtual utility functions are identical across bidders: $V_s =: V$ for all bidders s. As V is assumed to be strictly increasing, $V(x_i) > V(x_j) \Leftrightarrow x_i > x_j$, i.e., comparing bids in virtual utilities is equivalent to comparing bids in their actual values. Thus, the optimal mechanism can be simply interpreted as a sequence of second-price auctions among the bidders invited so far subject to period-specific reserve prices. By (21), in each period t, the reserves r_s^t are identical for all invited bidders $s \in \{1, \ldots, t\}$; furthermore, this uniform reserve price declines over time, as v_{t+1}^* does so. Section 3.3 presents a simple example of such a sequence of second-price auctions.

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 [10]. 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 in their setting and the optimal mechanism is simply a sequence of standard second-price auctions with a constant reserve price, with one more bidder invited in every period.

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

In this subsection we examine a specific example with two potential bidders that illustrates how the optimal mechanism can be computed. Suppose that x_i is uniformly distributed on [0, 1] for each potential bidder i, and 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

$$V_i(x) = 2x_i - 1. \tag{22}$$

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

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

$$\int_{(1+v_i^*)/2}^{1} \left[2x - 1 - v_i^*\right] dx = c_i \iff v_i^* = 1 - 2\sqrt{c_i}.$$

Since $c_1 < c_2 < \frac{1}{4}$, we get $v_1^* > v_2^* > 0$, 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 $V_1(x_1) \geqslant v_2^*$, i.e., if $x_1 \geqslant 1 - \sqrt{c_2}$. Otherwise, the mechanism continues to period 2, and bidder 1 wins if and only if $x_1 > \max\left\{x_2, \frac{1}{2}\right\}$. Since $x_2 \sim U[0, 1]$, the probability that bidder 1 wins is 0 if $x_1 < \frac{1}{2}$, x_1 if $\frac{1}{2} \leqslant x_1 < 1 - \sqrt{c_2}$, and 1 if $x_1 \geqslant 1 - \sqrt{c_2}$. Using (9), the expected payment of bidder 1 is

$$p_1(x_1) = \begin{cases} 0 & \text{if } x_1 < \dfrac{1}{2}, \\ x_1^2 - \int_{1/2}^{x_1} z\,dz = \dfrac{x_1^2}{2} + \dfrac{1}{8} & \text{if } \dfrac{1}{2} \leqslant x_1 < 1 - \sqrt{c_2}, \\ x_1 - \int_{1/2}^{1-\sqrt{c_2}} z\,dz - \int_{1-\sqrt{c_2}}^{x_1} dz = \dfrac{1 - c_2}{2} + \dfrac{1}{8} & \text{if } x_1 \geqslant 1 - \sqrt{c_2}. \end{cases} \tag{23}$$

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

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

The optimal mechanism can be implemented with the following sequence of second-price auctions, which is a simplest special case covered by our characterization in Section 3.2. Bidder 1 is offered the good at price $(1 - c_2)/2 + \frac{1}{8}$ in period 1 (a degenerate second-price auction with one participant); if he rejects the offer, he and bidder 2 participate in period 2 in a second-price auction with a reserve price equal to $\frac{1}{2}$. (Note, since $c_2 < \frac{1}{4}$, we have $(1 - c_2)/2 + \frac{1}{8} > \frac{1}{2}$, hence the reserve price declines over time, as predicted in Section 3.2.) Since both bidders bid their values in the period-2 second-price auction, bidder 2 wins if and only if $x_2 > \max\{x_1, \frac{1}{2}\}$ and his payment if he wins is $\max\{x_1, \frac{1}{2}\}$. 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 $\frac{1}{2}$ if he wins. Thus, if $x_1 \leqslant \frac{1}{2}$, 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

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

This expression is equal to the second line in Eq. (23).

Finally, to verify that bidder 1 accepts the offer in period 1 if and only if $x_1 \geqslant 1 - \sqrt{c_2}$, 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

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

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

In the general case with parallel search, it is difficult to identify the consequences of a change in search cost, except for the trivial fact that an increase in search costs decreases the seller's profits. In the current example however it is easy to identify the effects: an increase in c_1 (assuming that c_1 is still below c_2) does not affect the bidders' utilities but does decrease the seller's profit. On the other hand, an increase in c_2 lowers the reserve price in period 1 and hence makes it more likely that the mechanism will end in period 1. This in turn increases bidder 1's utility while decreasing bidder 2's utility.

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 Section 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 Section 4.2, we show an opposite effect: the optimal mechanism gives the seller an excessive incentive to search relative to the efficient mechanism. In Section 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.[12] Given these assumptions, the seller's optimal search mechanism is characterized in Section 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

$$\mathbb{E}_{x_i}\left[x_i - x_i^*\right]^+ = c_i. \tag{24}$$

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 PBE.[13]

4.1. Fewer participants

Since 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 exclude a bidder even before the search begins, even though that bidder has a positive probability of participation in an efficient mechanism. This is 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 $z \leqslant \overline{x}_i$, $\mathrm{E}_{x_i}[V_i(x_i) - z]^+$ and $\mathrm{E}_{x_i}[x_i - z]^+$ are strictly decreasing functions of z and $\mathrm{E}_{x_i}[V_i(x_i) - z]^+ \leqslant \mathbb{E}_{x_i}[x_i - z]^+$. 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 $v_i^* \geqslant x_0$. □

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

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

[13] 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.

to the previous one, as the lower fallback value makes it more attractive to continue the search. A simple case for this effect is when 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, $V'(x) \geqslant 1$. 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 than, 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, $x_i^* > x_j^* \iff c_i < c_j \iff v_i^* > v_j^*$. Thus, we can relabel the bidders so that $v_1^* \geqslant v_2^* \geqslant \cdots \geqslant v_n^*$ and $x_1^* \geqslant x_2^* \geqslant \cdots \geqslant x_n^*$. 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

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

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 $\alpha_j \in [0, 1)$. 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 Sections 4.1–4.3, a distortion due to asymmetry of information.

Proposition 3. *Assume $V_i \geqslant 0$ and $\underline{x}_i \geqslant 0$ 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) in the dynamic programming problem of optimal search,

$$\mathbb{E}_{x_{-J}} V_i(x_J, x_{-J}) > \mathbb{E}_{x_{-J}} V_j(x_J, x_{-J}) \iff W_i(x_i) > W_j(x_j). \tag{25}$$

Since $V_i \geqslant 0$ by assumption, the optimal mechanism never results in no sale. Hence (25) implies that the search procedure is equivalent to a standard Weitzman search 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 in α_i (since $x_i \geqslant 0$ 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 "influential" 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.

Acknowledgments

We are grateful for comments received from Dirk Bergemann, Eddie Dekel, Rich McLean, Andrew McLennan, Josh Pinkston, Michael Schwarz, Asher Wolinsky, Lixin Ye, an associate editor, two referees, and the seminar participants of "Endogenous Information Acquisition" in the 2006 Annual ASSA Meeting. Yossi Spiegel gratefully acknowledges the financial support of the IIBR and Charles Zheng thanks the NSF for Grant SES-0214471.

Appendix. A formal definition of revelation search mechanisms

Here we complement the informal definition of sequential mechanisms in Section 2.3 by a formal definition. This enables us to provide a formal link between Section 2.3 and the definitions introduced in Section 2.4. To this end we first define a decision tree for the search problem. Then we define a search procedure as a transition function on the decision tree, and payment scheme and disclosure policy as plans contingent on the nodes.[14] In the following, a decision node is called *principal node* if it is the principal's turn to move, and *agent node* if it is agents' turn to move.

The initial node of the decision tree, d^0, is a principal 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 $E^1 = \emptyset$, then the seller keeps the good without conducting any search. Each (d^0, E^1) is an *agent node*. The immediate

[14] This type of formalism was first developed by Zheng [20].

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 a profile of types that period-1 entrants reported as their realized types. Each (d^0, E^1, x^1) is a principal node at period 1.

Let $d^t := (d^0, E^1, x^1, \ldots, E^t, x^t)$ be a period-t principal node. The immediate successors of d^t are either (d^t, w) or (d^t, E^{t+1}), where $w \in \bigcup_{s=1}^t E^s \cup \{\text{seller}\}$ is the winner of the good (which could be the seller), and $E^{t+1} \subseteq I \setminus \bigcup_{s=1}^t E^s$, with $E^{t+1} \neq \emptyset$, is the set of period-$(t+1)$ entrants.[15] Each (d^t, w) is a *terminal node* specifying who gets the good, while each (d^t, E^{t+1}) is an agent node specifying the potential bidders who are invited to participate in period $t+1$. Note that if $I \setminus \bigcup_{s=1}^t E^s = \emptyset$, then the immediate successors of d^t are all terminal. The immediate successors of any agent 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 a profile of types that period-$(t+1)$ entrants reported as their realized types, and (d^t, E^{t+1}, x^{t+1}) is a period-$(t+1)$ principal node. Hence a decision tree is recursively defined.

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

$$(d^t)_+ := \left\{(d^t, w) : w \in \bigcup_{s=1}^{t} E^s \cup \{\text{seller}\}\right\} \cup \left\{(d^t, E^{t+1}) : \emptyset \neq E^{t+1} \subseteq I \Big\backslash \bigcup_{s=1}^{t} E^s\right\}$$

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

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

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 $(p_i^t)_{i \in \bigcup_{s=1}^t E^s} \in \mathbb{R}^{\bigcup_{s=1}^t E^s}$ 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 on the set of agent 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 type, 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 agent nodes $(d^0, E^1, \ldots, E^s)$ such that $i \in E^s$, with $s = 1, 2, \ldots$ (so i cannot tell any two of these nodes apart); hence i does not even know which period he is in, nor 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 principal 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 $\sigma(d^t)$.

[15] The assumption $E^{t+1} \neq \emptyset$ 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 *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 by $w(\tilde{x})$, 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:

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

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 because every player is risk neutral and types are independent.

References

[1] D. Bergemann, M. Pesendorfer, Information structures in optimal auctions, Cowles Foundation Discussion Paper 1323, August 2001.
[2] D. Bergemann, J. Välimäki, Information acquisition and efficient mechanism design, Econometrica 70 (2002) 1007–1033.
[3] D. Bergemann, J. Välimäki, Information in mechanism design, Cowles Foundation Discussion Paper 1532, August 2005.
[4] R. Burguet, Optimal repeated purchases when sellers are learning about costs, J. Econ. Theory 68 (1996) 440–455.
[5] O. Compte, P. 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] J. Crémer, Y. Spiegel, C.Z. Zheng, Optimal selling mechanisms with costly information acquisition, Working Paper, Institut d'Economie Industrielle (IDEI), Toulouse, August 12, 2003.
[7] S. Gal, M. Landsberger, A. Nemirovski, Participation in auctions, Mimeo, University of Haifa, Haifa, Israel, 2003.
[8] A. Gershkov, B. Szentes, Optimal voting schemes with costly information acquisition, Mimeo, Hebrew University of Jerusalem, University of Chicago, 2004.
[9] D. Levin, J.L. Smith, Equilibrium in auctions with entry, Amer. Econ. Rev. 84 (1994) 585–599.
[10] R.P. McAfee, J. McMillan, Search mechanisms, J. Econ. Theory 44 (1988) 99–123.
[11] P. Milgrom, I. Segal, Envelope theorems for arbitrary choice sets, Econometrica 70 (2002) 583–601.
[12] R.B. Myerson, Optimal auction design, Math. Methods Operations Res. 6 (1981) 58–73.
[13] W. Pesendorfer, A. Wolinsky, Second opinion and price competition: inefficiency in the market for expert advice, Rev. Econ. Stud. 70 (2003) 417–437.

[14] L. Rezende, Mid-auction information acquisition, Mimeo, Department of Economics, Stanford University, October 2000.
[15] M. Stegeman, Participation costs and efficient auctions, J. Econ. Theory 71 (1996) 228–259.
[16] T. Vishwanath, Parallel search for the best alternative, Econ. Theory 2 (1992) 495–507.
[17] M.L. Weitzman, Optimal search for the best alternative, Econometrica 47 (1979) 641–654.
[18] A. Wolinsky, Competition in a market for informed experts' services, RAND J. Econ. 24 (1993) 380–398.
[19] L. Ye, Optimal auctions with endogenous entry, Contrib. Theoretical Econ. 4 (2004) Article 8.
[20] C.Z. Zheng, VCG mechanisms in dynamic processes, Mimeo, ⟨http://faculty-web.at.northwestern.edu/economics/zheng/research/working/vcg/⟩, October 2005.