Optimal Search Auctions*

Jacques Crémer[†] Yossi Spiegel[‡] Charles Z. Zheng[§]

First draft: September 24, 2002
This draft: February 1, 2004

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. When bidders' types are independent (with possibly bidder-specific distributions) and their valuations are possibly interdependent, the seller's problem can be reduced to a (stochastic, history-contingent) 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. When bidders' types are correlated, the seller cannot fully extract the social surplus for sure but can extract it with an arbitrarily high probability.

*This version supersedes a draft titled "Optimal Auctions from a Search-Theoretic Viewpoint." For comments and discussions, we thank Eddie Dekel, Rich McLean, Andrew McLennan, Josh Pinkston, Michael Scwatz, and Asher Wolinsky. Zheng thanks the NSF for grant SES-0214471.

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

[‡]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

[§]Department of Economics, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208, Tel: 847-491-8213, email: czheng@northwestern.edu, http://faculty-web.at.northwestern.edu/economics/zheng/.

1 Introduction

Almost all the auction literature assumes that the set of bidders is either exogenous or determined in advance before the auction begins. But auctions based on this assumption are in general suboptimal, both socially and from the seller's viewpoint, if the bidders' participation in the auction is costly. In this paper we study the design of optimal (profit maximizing) auctions under the assumption that the seller needs to incur costs in order to contact prospective bidders and inform them about the auction. We show that the presence of these costs raises interesting dynamic questions such as how to sequence the order in which bidders are contacted and when to stop the auction process.

We start with an auction environment of one good, independent types with bidder-specific distributions, and interdependent ex post values across bidders. Initially, a prospective bidder is not aware of the seller's intention to sell the good. To attract his attention and allow him to participate, the seller must contact a prospective bidder and provide him with all the necessary information truthfully. In learning this information, a prospective bidder also becomes privately informed of his valuation of the good, before he can contract with the seller. Such individual-based contacts are necessary because the information to be conveyed is too complex and too costly to be mass-broadcasted. (For example, a prospective bidder needs hands-on training of the auction rules.) Thus, to attract a bidder, the seller needs to incur a bidder-specific cost, called *search cost* in the sequel. Then it is generally not optimal to contact all bidders at once: for instance, if the valuation of an early bidder turns out to be sufficiently high it is optimal to end the mechanism and sell the good to that bidder. Hence the seller designs a *search mechanism* that, contingent on history, specifies the order in which prospective bidders are contacted and invited to participate, the time at which the process ends, and the payments made by the bidders who have entered the mechanism.

In Section 2, we formally define the notion of search mechanism. Following some preliminary derivations in Section 3.1, we prove one of our main results, Theorem 1, in Section 3.2: the seller's problem can be reduced to an operation research (stochastic, history-contingent) search problem in which the surplus is measured in terms of the bidder's virtual utilities minus search costs. The seller can at most extract the virtual utilities of the bidders she has contacted and invited to participate in the mechanism because these bidders are privately informed about their valuations before they can agree to participate (the participation constraints are interim). Since the seller searches for the bidder with the highest virtual (rather than actual) utility, the profit maximizing mechanism creates distortions. Section 3.3.1 studies three of these distortions: fewer participants, longer search conditional on the same set of participants, and inefficient sequence of entry. Section 3.3.2 presents another implication: the seller wants noninfluential bidders to enter the mechanism before influential bidders.[1]

In Section 4, we relax the independence assumption and study a private value model with correlated types. It is well known that, in the corresponding traditional model, the seller's optimal mechanism is a full extraction mechanism that implements the socially efficient allocation and gives the seller the entire ex ante social surplus. With search costs, in contrast, achieving full extraction is harder because a bidder may be able to exclude the entry of rivals by submitting a bid that terminates the search process. Hence the seller may be unable to induce truth-telling by Crémer-McLean lotteries that condition a bidder's payment on all his rivals' reports. This observation is formalized by Proposition 4. Despite these difficulties, however, we construct in Theorem 2 a search mechanism that allows the seller to achieve full extraction with a probability arbitrarily close to one.

Our theorems extend existing results in traditional mechanism design theory by endogenizing the set of participants through a stochastic, history-contingent, search procedure.[2] In particular, McAfee and McMillan [6] have characterized profit maximizing search mechanisms but only considered the special case of independent private values and symmetric bidders. Burguet [1] has considered a similar private-value *i.i.d.*-bidder model, except that the participation constraint in his model is ex ante (before bidders become privately informed) instead of interim. In an earlier paper, Crémer, Spiegel, and Zheng [3], we studied

[1]At first glance, this result may appear to be at odd with the linkage principle. However, the linkage principle is not applicable here because of the sequential structure here. If the influential bidder enters first and buys the good right away, he faces little competition due to the absence of his rival; if the influential bidder does not buy the good right away, subsequent bidders would interpret this fact as a bad signal about the value of the good.

[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 are not aware of any work in that literature considering this problem.

the design of optimal selling mechanisms when potential bidders do not know their valuations at the outset but can learn them at a cost. A crucial assumption in that paper was that the seller can costlessly contact the bidders before they learn their valuations and offer them contracts. Hence the bidders' participation constraints were ex ante rather than interim as in the current paper. As a consequence, the profit maximizing search mechanism did not introduce any distortions compared to the socially optimal mechanism.[3]

2 The Model

2.1 Search costs

A seller has an indivisible good that can be sold to one bidder out of a finite set I of prospective bidders. Initially, none of the bidders is aware of the seller's intention to sell the good, the rules in the seller's auction, or the environment (who the other bidders are, how their valuations are distributed, etc.). In order to bring these information to a bidder i's attention, the seller needs to incur a commonly known bidder-specific fixed cost $c_i > 0$, to which we refer as *search cost*. This cost represents the cost of contacting bidder i and providing him with all the necessary information about the good, the mechanism, and the environment. This communication needs to be on a one-to-one basis because the information to be conveyed is too complex and too costly to be advertised to the whole world. We assume that the seller cannot lie about such information. In learning these information, the bidder also privately learns his ex post valuation for the good, before he can contract with the seller. (Hence bidders' participation constraints are interim.) If a bidder agrees to participate, he and the seller sign a binding contingency contract. In any period, bidders who have agreed in previous periods to participate are called *incumbents*. At any time, the

[3] Other authors have studied the optimal choice of auction formats with costly information acquisition, including Levin and Smith [5],Ye [13], Bergemann and Pesendorfer (2001), and Bergemann and Välimäki (2002) (the latter is a general mechanism design problem). In all of these papers however the acquisition of information is done by all agents before they participate and all agents participate in the mechanism simultaneously. Hence, none of these paper consider search mechanisms as in our paper. There are also papers that study the consequences of participation costs that to some extent correspond to the search costs in our model. See for example, Stegeman [10] and in Gal, Landsberger, and Nemirovski [4].

set of incumbents, their actions and messages up to that moment, and the set of *entrants* in the current period are assumed commonly known (among all the contacted bidders and the seller). Nonparticipants get zero payoff.

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, such that X_i is an interval of real numbers with infimum $\underline{x}_i$ and $f_i > 0$ over its interior. Types are independent across i. A vector of types $x := (x_i)_{i \in I} \in \times_{i \in I} X_i$ is called a *realized state*. As in Myerson [8], given any realized state 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 valuation. 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 $s \leq t'$ 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', and $-\sum_{t=s}^{\infty} \delta^{t-s} p_i^t$ if he does not get the good.

2.3 Search mechanisms

With search costs, 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 incumbents' messages, specifies whether she should stop the mechanism and keep the good or allocate it to one of the incumbent bidders, or whether she should continue and invite new bidders. Coupled with a payment scheme, we refer to such a contingent plan as the seller's *search mechanism*. Note that parallel search is allowed since there can be several entrants in any given period.[4]

[4] In the traditional mechanism design framework with zero search cost, search mechanisms are available for the seller but do not generate more revenues than mechanisms in which every bidder participates. In our model, search mechanisms can be better because they economize on the search costs.

A search mechanism works as follows. It defines a set ψ^1 of entrants in period 1 who send a vector of messages y_{ψ^1}. Then the set of period-2 entrants is a function $\psi^2(y_{\psi^1})$ of y_{ψ^1}. If the search stops at the end of period 2, the winner is determined by the messages from the set of incumbents, $I^2 := \psi^1 \cup \psi^2(y_{\psi^1})$. Let y_{I^2} be the profile of messages from the incumbents. Then a lottery $q(y_{I^2}) := (q_i(y_{I^2}))_i$ picks the seller or one of the incumbent as the final owner of the good, where $q_i(y_{I^2})$ is the probability that i is picked. Using i_* to denote the final owner of the good, each bidder i makes a payment $p_i^2(y_{I^2}, i_*)$. If the search continues, a new set $\psi^3(y_{\psi^2})$ of entrants is invited and the process continues.

Next we turn to a formal definition of search mechanisms. To make the definition succinct, we extend the functions ψ^t (that specifies who should enter in period t), $p_i^t(\cdot, i_*)$ (that specifies how much i should pay in period t in case i_* wins), and q (that specifies how to select the final owner) into functions of the entire profile y of messages from all bidders (subject to the conditions specified below):

$$\psi^t : Y \to 2^I, \qquad q : Y \to \Delta(I \cup \{0\}), \qquad p_i^t : Y \times (I \cup \{0\}) \to R,$$

where $Y := \times_{i \in I} Y_i$ with Y_i being the message space for bidder i, 2^I is the set of all subsets of I, and $\Delta(I \cup \{0\})$ is the set of lotteries that pick a final owner of the good from the set of bidders and the seller. From the functions ψ^t we define for all $y \in Y$

$$I^t(y) := \cup_{s=1}^t \psi^s(y),$$

which records the set of incumbents in each period, and we define for all $y \in Y$

$$\tau(y) := \max\{s = 1, 2, \ldots : \psi^s(y) \neq \emptyset\},$$

the period at which the search ends. For any $J \subset I$ and any $y \in Y$, let $y_J := (y_i)_{i \in J}$ and $y_{-J} := (y_i)_{i \notin J}$. Then a search mechanism corresponds to a list $((\psi^t, (p_i^t)_{i \in I})_{t=1}^{\infty}, q)$ such that

1. ψ^1 is constant on Y.
2. For any $t = 1, 2, \ldots$ and for any $y, y' \in Y$, if $I^t(y) = I^t(y')$ and $y_{I^t(y)} = y'_{I^t(y')}$, then
   a. $\psi^{t+1}(y) = \psi^{t+1}(y') \subseteq I \setminus I^t(y)$, and
   b. $q(y) = q(y')$ if $t = \tau(y)$

3. For all $y \in Y$, $q_i(y) = 0$ if $i \notin I^{\tau(y)}(y)$, and $p_i^t(y, \cdot) = 0$ if $i \notin I^t(y)$ or $t > \tau(y)$.

To calculate expected values, we require that these functions be measurable relative to the message space.

2.4 Revelation search mechanisms and search procedures

A *revelation search mechanism* is a search mechanism in which each bidder i's message space Y_i is i's type space X_i. The sequence $((\psi^t)_{t=1}^{\infty}, q)$ in a revelation search mechanism is called a *search procedure*.

Once selected, a search mechanism induces a multistage game. If this induced game has a perfect Bayesian equilibrium (PBE), the mechanism is said to be *equilibrium feasible*. A search mechanism is (seller-)*optimal* if it maximizes the seller's expected value of profits among all equilibrium feasible search mechanisms. A revelation search mechanism is said to be *incentive feasible* if the induced multistage game has a PBE where every invited bidder participates and is truthful. The next lemma is analogous to the revelation principle, and its proof is straightforward.

Lemma 2.1 (Revelation Principle for Search Mechanisms) *For any equilibrium feasible search mechanism, the seller can use an incentive feasible revelation search mechanism that replicates the equilibrium outcome of the former mechanism.*

2.5 Relation to optimal search theory

In traditional search theory there is no *asymmetric* information once the search cost has been incurred. 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 must design a search mechanism that induces the bidders to reveal their private information.

For the moment, suppose that after a bidder is contacted by the seller, his type becomes common knowledge. Then, the seller's *symmetric-information search problem* would be as follows: the seller has an *initial fallback reward* r_0 if she keeps the good. Given any realized state x, the seller's ex post *reward* from selling the good to participant i is $r_i(x)$ for some known $r_i : Y \to R$. The seller also bears the search costs. Hence, if the seller follows a search procedure $((\psi^t)_{t=1}^{\infty}, q)$ and given the realized state x the search stops at the end of period $\tau(x)$, the seller's expected value of net profit from the viewpoint of period 1 is

$$\mathrm{E}_x \left[\delta^{\tau(x)-1} \left[\sum_{i \in I} q_i(x) \left(r_i(x) - r_0 \right) \right] - \sum_{t=1}^{\infty} \delta^{t-1} \sum_{i \in \psi^t(x)} c_i \right], \tag{1}$$

where $q_i(x_i) = 0$ if $i \notin I^{\tau(x)}(x)$ as part of the definition of search mechanism (Item 3 in §2.3) and where E_x denotes the expected-value operator for functions of the random vector x. (The notations E_{x_i} and E_{x_J} in the sequel are analogous.)

A search procedure $((\psi^t)_{t=1}^{\infty}, q)$ is said to be *symmetric-information efficient* relative to the reward structure $(r_0, (r_i)_{i \in I})$ if the expression in (1) is maximized over all search procedures. A search procedure is said to be *socially efficient* if it is symmetric-information efficient relative to $(r_0, (r_i)_{i \in I})$ such that $r_0 = x_0$ and $r_i(x) = x_i + \sum_{j \in I \setminus i} e_{ij}(x_j)$ for all $i \in I$ and all realized state x.

3 Independent Types

3.1 Preliminary analysis

By Lemma 2.1, we can confine attention to revelation search mechanisms. Let $((\psi^t, (p_i^t)_{i \in I})_{t=1}^{\infty}, q)$ be such a mechanism. Let bidder i enter the mechanism in period t. Before submitting a report, he already knows the set J of incumbents and the profile x_J of their reported types (which is null if $J = \emptyset$), but he is still uncertain about the types $x_{-(J \cup i)}$ of the other bidders. Let i's report be $\hat{x}_i$. From the viewpoint of period t and conditional on $(x_J, \hat{x}_i)$, the discounted expected value of the winning probability of bidder $k \in I$ is

$$Q_k(\hat{x}_i \mid x_J) = \mathrm{E}_{x_{-(J \cup i)}} q_k(\hat{x}_i, x_J, x_{-(J \cup i)}) \delta^{\tau(\hat{x}_i, x_J, x_{-(J \cup i)}) - t}, \tag{2}$$

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

$$e_{-i}(\hat{x}_i \mid x_J) = \mathrm{E}_{x_{-(J\cup i)}}\left[\delta^{\tau(\hat{x},x_J,x_{-(J\cup i)})-t} q_i(\hat{x}, x_J, x_{-(J\cup i)}) \sum_{j\in I\setminus i} e_{ij}(x_j)\right]. \tag{3}$$

Analogously, one can calculate the discounted expected value of bidder i's total payment from the viewpoint of period t and conditional on $(x_J, \hat{x}_i)$. Denote this discounted expected value of payment by $P_i(\hat{x}_i \mid x_J)$. If bidder i's realized type is x_i, then his discounted expected utility from reporting $\hat{x}_i$ is

$$u_i(\hat{x}_i \mid x_i, x_J) = x_i Q_i(\hat{x}_i \mid x_J) + e_{-i}(\hat{x}_i \mid x_J) - P_i(\hat{x}_i \mid x_J) \tag{4}$$

from the viewpoint of period t. Given the independence of bidders' types, $Q_i(\hat{x}_i \mid x_J)$, $P_i(\hat{x}_i \mid x_J)$, and $e_{-i}(\hat{x}_i \mid x_J)$ are all independent of 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.

Proved by small extension of the routines in optimal auction theory, the next lemma says that the seller's problem is the same as a symmetric-information search problem with a distorted reward structure if the solution of this search problem happens to satisfy a monotonicity condition. In this distorted reward structure, the seller's fallback value is x_0 and her ex post gross reward from selling the good to bidder i is bidder i's virtual utility V_i, defined in the following for every realized state x:

$$V_i(x) := x_i - \frac{1 - F_i(x_i)}{f_i(x_i)} + \sum_{j\in I\setminus i} e_{ij}(x_j). \tag{5}$$

Lemma 3.1 *If a search procedure $((\psi^t)_{t=1}^{\infty}, q)$ is symmetric-information efficient relative to the reward structure $(x_0, (V_i)_{i\in I})$ and if*

$$\textit{the function } Q_i(\cdot \mid x_J) \textit{ is monotone nondecreasing, for all } i,\ J,\ \textit{and } x_J, \tag{6}$$

then there exists a payment scheme with which $((\psi^t)_{t=1}^{\infty}, q)$ constitutes a seller-optimal search mechanism.

Proof: By standard techniques (e.g., Myerson [8, Lemma 2]), the quasilinear form of (4) implies that the seller's problem is equivalent to maximizing the expected value of her profit among all revelation search mechanisms $((\psi^t, (p_i^t)_{i\in I})_{t=1}^{\infty}, q)$ subject to (6) and the following two constraints for any x_i and x_J:

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

$$u_i(\underline{x}_i \mid \underline{x}_i, x_J) \geq 0. \tag{8}$$

The solution of the seller's problem is unchanged when (7) is replaced by

$$u_i(\underline{x}_i \mid \underline{x}_i, x_J) = 0, \tag{9}$$

for if (9) did not hold, slightly raising the payments of all types would raise profits.

For any realized state x and any bidder $i \in I$, let t_i be the period in which the mechanism asks bidder i to enter. At period t_i, the realized reports x_J of the incumbents J are commonly known and are assumed to be truthful. By Eqs. (4), (8), and (9), the seller's expected net profit extracted from bidder i, viewed from period t_i, is

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

This, again by a standard argument (e.g., Myerson [8, Lemma 3]), is equal to

$$\mathrm{E}_{x_{-J}}\left[\delta^{\tau(x_J,x_{-J})-t_i} q_i(x_J, x_{-J})\left(V_i(x_J, x_{-J}) - x_0\right) - c_i\right],$$

where we have used Eqs. (2), (3), and (5). Viewed from period 1, the period $t_i(x)$ at which bidder i enters the mechanism is a random variable that depends on the realized state x. Thus, viewed from period 1, the seller's expected profit extracted from bidder i is

$$\begin{aligned} & \mathrm{E}_x\left[\mathrm{E}_{x_{-J}}\left[\delta^{\tau(x_J,x_{-J})-1} q_i(x_J, x_{-J})\left(V_i(x_J, x_{-J}) - x_0\right) - c_i\delta^{t_i(x)-1}\right]\right] \\ = \; & \mathrm{E}_x\left[\delta^{\tau(x)-1} q_i(x)\left(V_i(x) - x_0\right) - c_i\delta^{t_i(x)-1}\right]. \end{aligned} \tag{10}$$

Summing (10) over all $i \in I$, the seller's expected profit is equal to

$$\begin{aligned} & \mathrm{E}_x\left[\delta^{\tau(x)-1}\sum_{i\in I} q_i(x)\left(V_i(x) - x_0\right)\right] - \sum_{i\in I}\mathrm{E}_x\left[c_i\delta^{t_i(x)-1}\right] \\ = \; & \mathrm{E}_x\left[\delta^{\tau(x)-1}\sum_{i\in I} q_i(x)\left(V_i(x) - x_0\right) - \sum_{t=1}^{\infty}\delta^{t-1}\sum_{i\in\psi^t(x)} c_i\right], \end{aligned} \tag{11}$$

where the equality follows because $t_i(x)$ is the period at which bidder i enters the mechanism while $\psi^t(x)$ is the set of entrants at period t. The second line in (11) is just equal to the objective function (1) in the symmetric-information search problem with $r_0 = x_0$ and $r_i = V_i$. Thus, if $((\psi^t)_{t=1}^{\infty}, q)$ solves that search problem and satisfies (6), then it is an optimum for the seller if there is an associated payment scheme that satisfies (8) and (9). To this end, construct the payment scheme by setting

$$P_i(x_i \mid x_J) = x_i Q_i(x_i \mid x_J) + e_{-i}(x_i \mid x_J) - \int_{\underline{x}_i}^{x_i} Q_i(z \mid x_J)dz, \quad \text{for all } i \in I \setminus J. \tag{12}$$

Then Eq. (4) implies Eqs. (8) and (9). ■

In the absence of search costs, proving that the seller's objective is equal to (11) would have given us the traditional recipe of optimal auction: for almost every realized 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.

3.2 An optimal search mechanism

We have seen from Lemma 3.1 that the symmetric-information efficient search procedure, with rewards distorted into virtual utilities, is seller-optimal if it satisfies the monotonicity condition (6). Verifying (6) is nontrivial because the search literature does not seem to have considered the family of search problems encountered here.[5] Here we prove that any symmetric-information efficient search procedure satisfies (6).[6] That implies our main theorem. We need Assumption 1 which in the private-value case when e_{ij} are constants, is implied by the standard monotone hazard rate assumption. With interdependent values, it provides a virtual-utility analog of the single crossing property.

[5] Weitzman [12] and Vishwanath [11] considered only private values, while we allow interdependent values. But even for the private-value case, we are not aware of any general characterization of efficient search procedures that allow multiple entrants per period.

[6] In the private-value case where the e_{ij}s are all equal to zero, the revealed-preference argument in Appendix C of Crémer, Spiegel, and Zheng [3] yields the proof. The proof here is slightly more complicated due to the interdependency of valuation across bidders.

Assumption 1 *For each* $i, j \in I$, $x_i - \frac{1-F_i(x_i)}{f_i(x_i)}$ *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 the following revelation search mechanism is seller-optimal: its search procedure is symmetric-information efficient relative to the virtual-utility reward structure* $(x_0, (V_i)_{i\in I})$, *and its payment scheme satisfies (12) for all realized state* x.

This theorem follows from Lemma 3.1 and Lemma 3.3 which we prove below. To prove that lemma, we analyze the dynamic programming problem associated with a symmetric-information search problem relative to the reward structure $(x_0, (V_i)_{i\in I})$. At each period, the state variable (not to be confused with a realized state x) is (J, x_J), with $J \subseteq I$ being the set of incumbents and x_J the profile of their realized types. Denote $\pi(J, x_J)$ the optimal expected value of the seller's net reward relative to $(x_0, (V_i)_{i\in I})$ conditional on the state variable (J, x_J). The function π is defined recursively by the Bellman equation:

$$\pi(J; x_J) := \max\left\{x_0, \max_{j\in J} \mathrm{E}_{x_{-J}} V_j(x_J, x_{-J}), \delta \max_{K\subseteq I\setminus J}\left[\mathrm{E}_{x_K}\pi(J\cup K; x_J, x_K) - \sum_{k\in K} c_k\right]\right\}. \tag{13}$$

On the right-hand side of Eq. (13), the first term (x_0) is the seller's reward from stopping the search and keeping the good, the second term is the expected value of her reward from stopping and selling the good to an incumbent with the highest reward, and the third term is the optimal value from continuing the search. The equation for the case where $J = I$, is implied by Eq. (13) given the convention that the maximum over the empty set is zero.

Obviously any solution to the dynamic programming problem (13) yields a search procedure, which, depending on the value of π, either stops and gives the good to the seller or to one of the incumbents or continues and invites more bidders. If $\mathrm{E}_{x_J}\pi(J; x_J) - \sum_{i\in J} c_i \leq r_0$ for all $J \subseteq I$, the seller should keep the good and not contact any bidder. By induction on the size of $I \setminus J$, one can easily prove the fact that, given any state variable (J, x_J), a search procedure is symmetric-information efficient relative to reward structure $(x_0, (V_i)_{i\in I})$ if and only if it solves equation (13).

In a symmetric-information efficient search procedure relative to $(x_0, (V_i)_{i\in I})$, how does a change in an incumbent i's realized type x_i affect his probability of winning the good? To answer this question, we prove the next lemma, where $\pi_+(J, x_J)$ denotes the expected value of the distorted reward from continuing the search:

$$\pi_+(J, x_J) := \delta \max_{K\subseteq I\setminus J} \left[\mathrm{E}_{x_K} \pi(J \cup K; x_J, x_K) - \sum_{k\in K} c_k \right]. \tag{14}$$

Lemma 3.2 *If $i \in J \subseteq I$ and x_J is a profile of realized types on J, then $\mathrm{E}_{x_{-J}} V_i(x_J, x_{-J})$ and $\pi_+(J, x_J)$ are absolutely continuous functions of x_i; whenever their derivatives exist,*

$$\frac{\partial}{\partial x_i} \mathrm{E}_{x_{-J}} V_i(x_J, x_{-J}) > \frac{\partial}{\partial x_i} \max_{j\in J\setminus i} \mathrm{E}_{x_{-J}} V_j(x_J, x_{-J}) \quad \textit{and} \tag{15}$$

$$\frac{\partial}{\partial x_i} \mathrm{E}_{x_{-J}} V_i(x_J, x_{-J}) \geq \frac{\partial}{\partial x_i} \pi_+(J, x_J), \tag{16}$$

and the equality in (16) holds only if $\mathrm{E}_{x_{-J}} V_i(x_J, x_{-J}) \geq \pi_+(J, x_J)$.

Proof: Denote $H_i(x_i) := x_i - \frac{1-F_i(x_i)}{f_i(x_i)}$. By Eq. (5), $\frac{\partial}{\partial x_i} \mathrm{E}_{x_{-J}} V_i(x_J, x_{-J}) = H_i'(x_i)$, and, for all $j \neq i$, x_i enters $V_j(x)$ only through the term $e_{ji}(x_i)$. Hence Assumption 1 implies (15).

To prove (16), we use a revealed-preference argument. Let $((\psi^t)_{t=1}^\infty, q)$ denote a search procedure that solves Eq. (13). Given state variable (J, x_J), consider a deviant plan of replacing the x_i in x_J by an $\hat{x}_i \in X_i$ and carrying out subsequent search according to $((\psi^t)_{t=1}^\infty, q)$ with the revised state variable $(J, x_{J\setminus i}, \hat{x}_i)$. Let $\hat{\pi}_+(J, x_J; \hat{x}_i)$ be the expected value of virtual utility from this deviant plan, discounted back to the current period t. Then x_i enters $\hat{\pi}_+(J, x_J; \hat{x}_i)$ only in the term

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

Hence

$$\frac{\partial}{\partial x_i} \hat{\pi}_+(J, x_J; \hat{x}_i) = Q_i(\hat{x}_i \mid x_{J\setminus i}) H_i'(x_i) + \sum_{k\in I\setminus i} Q_k(\hat{x}_i \mid x_{J\setminus i}) e_{ki}'(x_i),$$

where $Q_k(\hat{x}_i \mid x_{J\setminus i})$ denotes the discounted expected value of bidder k's winning probability from the viewpoint of period t conditional on the profile $(x_{J\setminus i}, \hat{x}_i)$ (defined in Eq. (2), with J there replaced by $J \setminus i$ here). As $((\psi^t)_{t=1}^\infty, q)$ solves the dynamic programming problem given the state variable (J, x_J),

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

Thus, the Milgrom-Segal envelope theorem ([7]) implies that $\pi_+(J, x_J)$ is an absolutely continuous function of x_i and, whenever its derivative exists,

$$\frac{\partial}{\partial x_i}\pi_+(J, x_J) = Q_i(x_i \mid x_{J\setminus i})H_i'(x_i) + \sum_{k\in I\setminus i} Q_k(x_i \mid x_{J\setminus i})e_{ki}'(x_i). \tag{17}$$

Thus, Assumption 1 implies (16). If its equality holds, bidder i wins almost surely in subsequent search if search were to continue. Then continuing search is dominated by awarding the good to bidder i right now, due to search costs. Hence $\mathrm{E}_{x_{-J}}V_i(x_J, x_{-J}) \geq \pi_+(J, x_J)$. ■

Now we are ready to prove the lemma that immediately implies Theorem 1.

Lemma 3.3 *Given Assumption 1, if a search procedure is symmetric-information efficient relative to reward structure $(x_0, (V_i)_{i\in I})$, then it satisfies the monotonicity condition (6).*

Proof: Let $J \subseteq I$ be the set of incumbents and let $i \in J$. We shall prove (6) by induction on the size of $I \setminus J$. The case of $J = I$ follows directly from (15). 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 symmetric-information efficient search procedure solves the problem

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

Lemma 3.2 implies that the probability $g(x_i)$ for bidder i to win in the current period is monotone nondecreasing in x_i. The induction hypothesis implies that the probability $h(x_i)$ (discounted back to next period) that he wins later, conditional on the event that he does not win in the current period, is monotone nondecreasing in x_i. Thus, his total discounted winning probability $g(x_i) + (1 - g(x_i))\delta h(x_i)$ is monotone nondecreasing in x_i, as desired. ■

3.2.1 An optimal mechanism with private values

Let us illustrate Theorem 1 for a special private-value case where $e_{ij} = 0$ for all $i, j \in I$ and there is no discounting. With private values, a bidder's virtual utility becomes a function of only his own type (hence we write $V_i(x_i)$ instead of $V_i(x)$):

$$V_i(x_i) = x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}. \tag{18}$$

By Theorem 1, a solution to the dynamic programming problem (13), with the private-reward structure $(x_0, (V_i)_{i \in I})$, is optimal for the seller. The solution in such private-reward cases has been characterized by Weitzman [12]. (Weitzman allowed only one entrant per period, which we can do with no loss of generality when there is no discounting.) In our auction environment, the solution to the problem is: For every bidder, calculate a *cutoff* level such that if the seller's current fallback reward is below this cutoff then it is worthwhile to invite that bidder before ending the search. In any period t, if the seller's current fallback reward is greater than or equal to the cutoffs of all bidders who were not yet invited, the search terminates. Otherwise, the search continues to period $t+1$ and the seller contacts the bidder with the highest cutoff among all bidders who were not invited up to that point.

The cutoffs in the Weitzman algorithm, given a *private*-reward structure $(r_0, (r_i)_{i \in I})$, are computed as follows. Suppose that before the seller faces bidder i, she already has the opportunity to get a fallback reward z. The net change in the seller's expected payoff from inviting bidder i is equal to

$$\mathrm{E}_{x_i} \left[r_i(x_i) - z \right]^+ - c_i,$$

where

$$[y]^+ := \max\{y, 0\}.$$

This expression reflects the fact that if $r_i(x_i) < z$, then the seller's fallback reward remains equal to z but if $r_i(x_i) > z$ then $r_i(x_i)$ becomes the new fallback reward. The seller gains from inviting bidder i if and only if the resulting net change in her expected payoff is positive. Hence the cutoff for bidder i is equal to the solution r_i^* for the following equation (the existence and uniqueness of the solution are obvious):

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

When the reward structure is distorted into virtual utilities ($r_0 = x_0$ and $r_i = V_i$ for all $i \in I$), the optimal search procedure follows Weitzman's algorithm with the cutoffs, called *seller-optimal cutoffs* v_i^*, implicitly defined by

$$\mathrm{E}_{x_i} \left[V_i(x_i) - v_i^* \right]^+ = c_i. \tag{19}$$

That yields the following optimal mechanism: Relabel the bidders if necessary so that $v_1^* \geq \cdots \geq v_n^*$ ($n = \#I$). If $0 \geq v_1^*$, the seller keeps the good. Otherwise, invite bidder 1. If the search continues to period t and $(\hat{x}_s)_{s=1}^{t-1}$ is the profile of reports so far, invite bidder t; based on his report $\hat{x}_t$, set the bidder's payment such that its expected value is equal to the amount $P_t(\hat{x}_t \mid (\hat{x}_s)_{s=1}^{t-1})$ specified by Eq. (12). If $v_{t+1}^* > \max_{s\leq t} V_s(\hat{x}_s))$, continue the search and invite bidder $t+1$. Otherwise or if all bidders have been invited, stop the search and give the good to the participant whose reported virtual utility is at least x_0 and is highest among all participants. If no such participant exists, the seller consumes the good.[7]

From the above search procedure and Eq. (12), one can derive the following payment scheme for any bidder t who enters at period t. Since $v_{t+1}^* \leq v_t^*$, either $v_{t+1}^* \leq \max_{s<t} V_s(\hat{x}_s) \leq v_t^*$ or $\max_{s<t} V_s(\hat{x}_s) < v_{t+1}^* \leq v_t^*$. ($\max_{s<t} V_s(\hat{x}_s) < v_t^*$ because otherwise search would have stopped before t.) In the first case, the search will stop before period $t+1$ regardless of bidder t's report, and bidder t's payment scheme is a take-it-or-leave offer at the price equal to the bidder's minimum realized value that allows him to outbid the incumbents in terms of virtual utilities. In the second case, bidder t is pivotal to the decision of search continuation. If his report exceeds $V_t^{-1}(v_{t+1}^*)$, he buys the good now and pays a price equal to $V_t^{-1}(v_{t+1}^*)$ minus a discount (the discount is needed to counterbalance his incentive of prolonging the search by underbidding); if his report is below $V_t^{-1}(v_{t+1}^*)$, search continues and bidder t is committed to a payment plan whose discounted expected value is determined by Eq. (12).

3.3 Policy implications

3.3.1 Distortion on social efficiency

In the traditional optimal auction theory, asymmetric information can lead to inefficiency 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. To focus on the effect of asymmetric information, we consider in this subsection only a case of private values where $e_{ij} = 0$ for all $i, j \in I$ and there is no discounting.

[7] The sequence of entry is predetermined in this case because values are private and there is no discounting. In general, the sequence of entry is stochastic.

With private values and no discounting, Weitzman's algorithm solves the symmetric-information search problem. When the reward structure is based on the actual utilities ($r_0 = x_0$ and $r_i(x_i) = x_i$ for all $i \in I$), the socially efficient search procedure follows Weitzman's algorithm with the cutoffs x_i^*, called *efficient cutoffs*, implicitly defined by the equation

$$\mathrm{E}_{x_i}\left[x_i - x_i^*\right]^+ = c_i. \tag{20}$$

Crémer, Spiegel, and Zheng [3] have proved that this procedure can always be implemented by a perfect Bayesian equilibrium (PBE).[8] Theorem 1 above implies that the seller-optimal search procedure, which follows Weitzman's algorithm with the cutoffs v_i^* defined by Eq. (19) replacing the efficient cutoffs $(x_i^*)_{i \in I}$, is also implemented by a PBE. In what follows we can compare the two search procedures.

Fewer Participants. Because a bidder's actual utility exceeds his virtual utility, the benefit of including a bidder is lower in a seller-optimal search mechanism than in a socially efficient search mechanism if the fallback rewards are the same. The optimal mechanism may completely exclude a bidder even before the search begins, while that bidder has a positive probability of participation in an efficient mechanism.

Proposition 1 *From the standpoint of period 1, every bidder i's probability of participation in a socially efficient mechanism is positive whenever his probability of participation in a seller-optimal mechanism is positive, but bidder i's probability of participation in a seller-optimal mechanism can be zero even when his probability of participation in a socially efficient mechanism is positive.*

Proof: Note that for all z, $\mathrm{E}_{x_i}\left[V_i(x_i) - z\right]^+ \leq \mathrm{E}_{x_i}\left[x_i - z\right]^+$, with strict inequality for all $z < \sup X_i$; also note

$$\frac{d}{dz}\left(\mathrm{E}_{x_i}\left[V_i(x_i) - z\right]^+\right) \leq 0, \qquad \frac{d}{dz}\left(\mathrm{E}_{x_i}\left[x_i - z\right]^+\right) \leq 0,$$

with strict inequalities for all $z < \sup X_i$. 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

[8] 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.

efficient mechanism if $x_i^* > x_0$ and a positive probability of participation in the seller-optimal mechanism *only* if $v_i^* \geq x_0$. ■

Longer Search. As a bidder's virtual utility is less than his value, the seller's fallback value in an optimal mechanism is less than her fallback value in an efficient mechanism. That leads to an effect opposite to the previous one. The lower fallback value makes it more attractive to continue the search. A simple case for this effect is that bidders' types are drawn from an identical distribution F with density f, so their virtual utility functions are the same, though their participation costs may be different. While the cost of an additional searching period is the same in both 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, which is measured in virtual utilities, is approximately $V'(x)\Delta x$. Under the monotone hazard rate assumption that $f(x_i)/(1 - F(x_i))$ is weakly increasing, $V'(x) \geq 1$. Hence, other things equal, the seller is more willing to continue searching than a benevolent social planner would.

Proposition 2 *Assume that types x_i are identically distributed across bidders, with V denoting the common virtual utility function and V' its derivative (though their participation costs c_i may be different), and assume that $V' > 1$ and $v_i^* > x_0$ for all $i \in I$. Then a seller-optimal search lasts at least as long as an efficient search, and with a positive probability the former lasts longer than the latter.*

Proof: First, note the fact that $V^{-1}(v_i^*) > x_i^*$ for all $i \in I$: Let

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

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

Second, the queue 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 $x_1^* \geq x_2^* \geq \cdots \geq x_n^*$. By the assumption $v_i^* > 0$, the optimal mechanism for sure conducts the search in period one. We shall claim that the seller-optimal search continues from period t to period $t+1$ in a higher probability than the efficient search procedure does from t to $t+1$. 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 $\max\{x_0, x_1, \ldots, x_t\} < x^*_{t+1}$; by the fact $V^{-1}(v^*_i) > x^*_i$ and the assumption $v^*_i > x_0$,

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

Hence the seller-optimal search continues to period $t+1$. Thus, the optimal procedure continues if 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 seller-optimal search procedure continues. This proves our claim. ■

Inefficient Queue of Entry. Determined by different sets of cutoffs, the queues of entry in the efficient mechanism and in the optimal mechanism can be different. Here is an example where the seller-optimal queue of entry is the reverse of the efficient queue: there are two bidders and the seller's value x_0 is zero; bidder 1's type is uniformly distributed on $[\underline{x}_1, \overline{x}_1]$; bidder 2's type is drawn from an exponential distribution $F_2(x_2) := 1 - \exp(-\lambda x_2)$. We calculate the virtual utility functions and cutoffs of the two bidders:

$$\begin{aligned} V_1(x_1) &= 2x_1 - \overline{x}_1; & V_2(x_2) &= x_2 - 1/\lambda; \\ x^*_1 &= \overline{x}_1 - \sqrt{2c_1(\overline{x}_1 - \underline{x}_1)}; & x^*_2 &= -\ln(\lambda c_2)/\lambda; \\ v^*_1 &= \overline{x}_1 - 2\sqrt{c_1(\overline{x}_1 - \underline{x}_1)}; & v^*_2 &= -\left(1 + \ln(\lambda c_2)\right)/\lambda. \end{aligned}$$

Since $x^*_i > v^*_i$, there exist two numbers a and b such that $v^*_1 < a < b < x^*_1$. Let

$$\lambda := 1/(b-a); \quad c_2 := \exp(-\lambda b)/\lambda.$$

The choice of λ and c_2 implies that $x^*_1 > x^*_2$ and $v^*_1 < v^*_2$. Thus, bidder 2 enters first in the optimal mechanism, whereas bidder 1 enters first in the efficient mechanism.

In this example, search costs can be arbitrarily small (when $c_1 \to 0$, $b - a \to 0+$ and $c_2 \to 0$). Thus, when search costs go to zero, the distortion due to asymmetric information persists, while the delay of market clearance in the traditional search models vanishes.

3.3.2 Delayed participation of influential bidders

When bidders' values are interdependent, we can address the following question: If bidder i has stronger influence on others than bidder j, should the seller let i enter before j or vice versa? For simplicity, assume 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 $i \neq j$. We can therefore regard bidders with higher α's as more influential.

By the Milgrom-Weber linkage principle, one might think that the seller would rather have influential bidders enter earlier. The principle, however, is not applicable here, because only bad signals can affect future potential entrants. Suppose that an influential bidder enters first. If he does not purchase the good, then later entrants will take this as a bad signal about the good and hence reduce their willingness-to-pay. If the influential bidder does purchase the good, then the potential entrants, having no chance to compete, do not contribute to the bidding competition. Thus, a seller would like noninfluential bidders to enter the mechanism before influential bidders do.

Proposition 3 *For every $i \in I$, assume: $V_i \geq 0$, $\underline{x}_i \geq 0$, and there is a number α_i such that $e_{ji}(x_i) = \alpha_i x_i$ for all x_i and all $j \neq i$. Also assume that $\delta = 1$ and search costs and type-distributions are identical across $i \in I$. Then the higher α_i, the later and less probable is i's entry in a seller-optimal 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),

$$\mathrm{E}_{x_{-J}} V_i(x_J, x_{-J}) > \mathrm{E}_{x_{-J}} V_j(x_J, x_{-J}) \Longleftrightarrow W_i(x_i) > W_j(x_j). \tag{21}$$

Since $V_i \geq 0$ by assumption, the seller-optimal search mechanism never results in no sale. Hence (21) implies that the search procedure is equivalent to a symmetric-information efficient procedure relative to the reward structure $(r_0, (r_i)_{i \in I})$ such that $r_0 = -\infty$ and $r_i(x) = W_i(x_i)$ for all i and all x. In this transformed search problem, the reward r_i is a function of only x_i. Hence the search procedure is obtained via Weitzman's algorithm, with the cutoffs w_i^* implicitly defined by the next equation (the existence and uniqueness of w_i^* are obvious):

$$\mathrm{E}_{x_i}\left[W_i(x_i) - w_i^*\right]^+ = c_i. \tag{22}$$

Thus, it suffices to prove that w_i^* is strictly decreasing in α_i for each $i \in I$. To this end, note that $\mathrm{E}_{x_i}\left[W_i(x_i) - w_i^*\right]^+$ is a strictly decreasing function of w_i^* and moreover note that $W_i(x_i)$ is a strictly decreasing function of α_i (since $x_i \geq 0$ by assumption). Hence (22) implies the desired assertion. ■

4 Correlated Types

Thus far, we have assumed that the bidders' types, $x_1, ..., x_n$, are stochastically independent across bidders. In this section we relax this assumption. It is well known that when bidders' types are correlated, the seller can generically use lotteries that condition the payment of each bidder on the reports of other bidders to fully extract the entire social surplus (see Crémer-McLean [2, Theorem 2]). The presence of search costs, however, renders such full extraction difficult because a search procedure continues or stops depending on incumbents' reports. Hence an incumbent can prevent the entry of rivals and therefore make it impossible to use the appropriate lotteries. We begin this section by analyzing the source of this difficulty and then prove that the seller can overcome this difficulty and achieve almost full extraction albeit this may require the use of arbitrarily large transfers.

We consider a correlated private value case with a finite nonsingleton set I of bidders. The seller's value is normalized to zero. We assume that bidder i's ex post utility from winning the good is equal to the realized value of his type x_i, and we assume that the set of possible realized values of x_i is a finite set X_i.

4.1 A case where full extraction is impossible

If the socially efficient search procedure invites a single bidder in period 1, this bidder may be able to make a report that induces the seller to stop the mechanism. Consequently, the seller is unable to use lotteries to fully extract the surplus from this bidder. The next proposition makes this statement precise.

Proposition 4 *If the socially efficient search procedure invites a single bidder in period 1, and if there are at least two possible realized values of his type at which the efficient procedure calls for awarding the good to the bidder without further search, then the seller cannot fully extract the maximum social surplus in the symmetric-information search problem.*

Proof: Consider any revelation search mechanism that implements the efficient search procedure. (There is no need to consider any other mechanisms by the revelation principle, Lemma 2.1.) Suppose the efficient procedure invites only bidder i in period 1 and awards the good to i if $x_i = L$ or $x_i = H$, with $L, H \in X_i$ and $L < H$. Then, if $x_i = L$ or $x_i = H$, all other bidders are nonparticipants and bidder i's payment cannot depend on their realized signals. Hence incentive compatibility requires that bidder i pay the same amount when $x_i = L$ or $x_i = H$. Moreover, interim participation constraint when $x_i = L$ implies that this payment cannot exceed L. Hence the seller cannot fully extract the social surplus. ■

The following example illustrates Proposition 4 and shows that the hypothesis in the proposition is nonvacuous. Suppose that there are two ex ante identical bidders, 1 and 2. The discount factor is one, and the search cost is $c > 0$. Each bidder's type can take three possible values: L, M, and H, with $L < M < H$. The joint probability of (x_1, x_2) is:

		x_2: L	M	H
x_1	L	f_{LL}	f_{LM}	f_{LH}
	M	f_{LM}	f_{MM}	f_{MH}
	H	f_{LH}	f_{MH}	f_{HH}

The entries in the matrix are probabilities that sum up to one. For instance, $\Pr(x_1 = M, x_2 = H) = f_{MH}$. To ensure that it is socially efficient to invite at least one bidder, assume

$$f_L L + f_M M + f_H H > c, \tag{23}$$

where $f_L := f_{LL} + f_{LM} + f_{LH}$, $f_M := f_{LM} + f_{MM} + f_{MH}$, and $f_H := f_{LH} + f_{MH} + f_{HH}$. Note that since there is no discounting, there is no loss of efficiency to invite only one of the ex ante identical bidders, say bidder 1, in period 1. If $x_1 = H$, then obviously search should stop and bidder 1 should get the good. Otherwise, if

$$\frac{f_{LM}M + f_{LH}H - (f_{LM} + f_{LH})L}{f_L} < c < \frac{f_{MH}(H - M)}{f_M}, \tag{24}$$

then it is optimal to stop if $x_1 = L$ but continue if $x_1 = M$.[9] Obviously the set of parameters satisfying both (23) and (24) is nonempty (also not nongeneric, as (23) and (24) are strict inequalities). Hence Proposition 4 implies that full extraction is impossible.

4.2 Almost full extraction

Although full extraction may be impossible, it is nonetheless possible to modify the efficient search procedure and achieve almost full extraction of surplus. This modification requires that the procedure always continue with a positive probability, thereby eliminating the first entrant's ability to exclude rivals.

Let us start with the previous example, assuming that both (23) and (24) are satisfied and assuming that the above joint probability matrix satisfies the cone condition for full extraction (Crémer and McLean [2, Theorem 2]; stated here as Assumption 2). Pick any small $\epsilon > 0$. If bidder 1 reports M, continue search as in the efficient procedure. If bidder 1 reports L or H, stop search as in the efficient procedure with probability $1-\epsilon$, and continue search with probability ϵ. Since the procedure reaches bidder 2 with a positive probability and since the Crémer-McLean cone condition is assumed to hold, it is possible to design a Crémer-McLean lottery for bidder 1 that induces him to make a truthful report in period one. It remains to show that it is also possible to design such a lottery for bidder 2. To this end, suppose that bidder 2 is unaware of bidder 1's report. If the mechanism reaches bidder 2, the posterior joint probability measure from bidder 2's viewpoint corresponds to the following matrix up to normalization:

[9] If $x_1 = L$, then stopping yields a social surplus L while contacting bidder 2 and allocating the good to the higher-value bidder yields a social surplus of $\frac{f_{LL}L+f_{LM}M+f_{LH}H}{f_L} - c$. Likewise, if $x_1 = M$, then stopping yields M, while contacting bidder 2 and allocating the good to the higher-value bidder yields $\frac{(f_{LM}+f_{MM})M+f_{MH}H}{f_M} - c$. When (24) holds, $L > \frac{f_{LL}L+f_{LM}M+f_{LH}H}{f_L} - c$ and $M < \frac{(f_{LM}+f_{MM})M+f_{MH}H}{f_M} - c$.

		x_2		
		L	M	H
	L	ϵf_{LL}	ϵf_{LM}	ϵf_{LH}
x_1	M	f_{LM}	f_{MM}	f_{MH}
	H	ϵf_{LH}	ϵf_{MH}	ϵf_{HH}

For instance, $\Pr(x_1 = L \mid x_2 = H) = \epsilon f_{LH}/(\epsilon f_{LH} + f_{MH} + \epsilon f_{HH})$. To see that the associated conditional probability matrix satisfies the Crémer-McLean cone condition, suppose by way of negation that one of the column vectors, say the first column, belongs to the closed cone generated by the other two column vectors. Then for some nonnegative numbers λ_1, λ_2:

$$\begin{aligned}
\epsilon f_{LL} &= \lambda_1 \epsilon f_{LM} + \lambda_2 \epsilon f_{LH}; \\
f_{LM} &= \lambda_1 f_{MM} + \lambda_2 f_{MH}; \\
\epsilon f_{LH} &= \lambda_1 \epsilon f_{MH} + \lambda_2 \epsilon f_{HH}.
\end{aligned}$$

But then the *prior* probability measure also violates the cone condition, a contradiction. Hence it is also possible to induce truth-telling from bidder 2 by offering him a Crémer-McLean lottery contingent on bidder 1's report.

The above mechanism implements the efficient search procedure with probability $1-\epsilon$, where ϵ can be arbitrarily small. With both bidders being truthful, the seller fully extracts the social surplus in the event that the efficient procedure is implemented. Hence she obtains the maximum social surplus with probability at least $1-\epsilon$.

To allow the above mechanism, we expand the definition of search procedure to include the possibility of randomization on whether to continue and on the set of new entrants to be invited. We also replace the assumption that previous messages are common knowledge (§2.1) by the assumption that an incumbent's message is unknown to everyone else unless revealed by the mechanism. Without this, a lottery contingent on incumbents' messages could be degenerate. To ensure the credibility of the seller's mechanism, we further assume that a mechanism, once selected, is operated by a neutral trustworthy mediator.[10]

[10] These changes do not alter the results in the case of independence. Randomization does not make the seller better-off because a tiny tremble does not lead to a discontinuous change in the information rent. Nor can she do better by hiding previous messages in the independence case: as long as the mechanism is

Let $f((x_i)_{i\in I})$ denote the joint prior probability of the realized state $(x_i)_{i\in I}$, with $x_i \in X_i$ for each i. Assume that $f((x_i)_{i\in I}) > 0$ for all realized states $(x_i)_{i\in I}$. Let $X_{-i} := \times_{j\neq i} X_j$. For any $x_{-i} \in X_{-i}$, let $f_{-i}(x_{-i} \mid x_i)$ denote the probability of x_{-i} being the profile of realized types of all bidders but i, conditional on i's type being x_i. As each possible state has a positive prior probability, $f_{-i}(x_{-i} \mid x_i)$ is well defined. Note that $f_{-i}(\cdot \mid x_i)$ is a vector whose length is equal to the size of X_{-i}. The next assumption is exactly the cone condition in Crémer and McLean [2, Theorem 2].

Assumption 2 *For any bidder i and for any $x_i \in X_i$, the vector $f_{-i}(\cdot \mid x_i)$ does not belong to the closure of the cone generated by the vectors in the family $\{f_{-i}(\cdot \mid x_i') : x_i' \in X_i \setminus \{x_i\}\}$.*

As in the previous example, the main idea in the proof of the next theorem is to ensure a positive probability for the event of full participation, in which case Crémer-McLean lotteries can be carried out. Although this probability may be tiny, the lotteries can be scaled up to deter lying. The only complication in the proof is due to the fact that entrants can learn from the history of entry. To achieve full extraction, the seller needs to ensure that every entrant's posterior belief will satisfy the condition for full extraction, which requires a bidder's posterior conditional probabilities to be well defined. To guarantee that, we generalize the above ϵ-deviation technique into totally mixed strategies at the end of every period so that a new entrant always assigns positive posterior probabilities to any possible realized state.

Theorem 2 *Given Assumption 2, for any $\eta > 0$ there exists a search mechanism with which the seller obtains the maximum social surplus of the symmetric-information search problem with a probability at least $1-\eta$.*

Proof: Pick a sufficiently small $\epsilon > 0$ such that $1-\eta < (1-\epsilon)^{n-1}$ (n being the size of I). Consider the following mechanism: In period one, invite the entrants prescribed by the efficient procedure. In every period t, offer a menu of Crémer-McLean lotteries (specified later) to every period-t entrant, then solicit secret reports from them. If all bidders have participated, stop. Otherwise, with probability $1-\epsilon$ follow the instruction of the efficient

operated by a neutral trustworthy mediator rather than the seller herself, the seller does not know more than the bidders about previous messages, and hence she cannot take advantage of the hidden history.

procedure in period $t+1$, and with probability ϵ randomly pick, with equal probability, a nonempty set of bidders who are not yet incumbents and invite them in period $t+1$. If search stops, sell the good to a highest-value participant at a price equal to his reported value. In addition, participants make transfers according to their Crémer-McLean lotteries.

If all participants are truthful, the efficient search procedure is implemented with probability at least $(1-\epsilon)^{n-1}$. If the lotteries have the Crémer-McLean property of ensuring zero expected payoff for truth-tellers and sufficiently large negative payoffs for liars, then participants are indeed truthful so the seller obtains the entire social surplus if the efficient procedure is implemented. By definition of efficiency, that surplus is the maximum surplus in the symmetric-information search problem. Thus, the proof is complete if such a Crémer-McLean lottery exists for every participant i.

To this end, consider any bidder i who enters at period $t = 1, 2, \ldots$. Given i's report, $\hat{x}_i$, suppose that he is offered the following lottery: if search ends before all potential bidders participate, bidder i gets zero payoff; otherwise (full participation) and if x_{-i} is the profile of reports from all potential bidders but i, then bidder i gets a payoff equal to $\gamma_i(\hat{x}_i)g_i(\hat{x}_i, x_{-i})$ for some functions γ_i and g_i. We shall prove that there exist such functions for which the lottery has the desired Crémer-McLean property. By the totally mixed strategy described above, given any profile x_{-i} there is a unique positive probability $a(x_{-i})$ with which the mechanism coupled with x_{-i} leads to the observed sequence of entry up to the current period. (Actually $a(x_{-i})$ depends only on the reports of the incumbents before i enters.) Derived from the design of the mechanism, $a(x_{-i})$ is commonly known. Likewise, given any x_{-i} and $s = 0, 1, 2, \ldots$, there is a unique positive probability $\beta(\hat{x}_i, x_{-i}, s)$ with which the mechanism coupled with $(\hat{x}_i, x_{-i})$ leads to the observed sequence of entry up to the current period and will end with full participation in period $t+s$. Given $(\hat{x}_i, x_{-i}, s)$, this probability is commonly known. Let bidder i's actual type be x_i. Denote

$$\begin{aligned} b(x_i) &:= \sum_{x'_{-i} \in X_{-i}} a(x'_{-i}) f_{-i}(x'_{-i} \mid x_i); \\ G_i(\hat{x}_i, x_{-i}) &:= \sum_{s=0}^{\infty} \delta^s \beta(\hat{x}_i, x_{-i}, s) g_i(\hat{x}_i, x_{-i}). \end{aligned}$$

Then bidder i's expected payoff from the lottery, viewed from the current period, is equal to

$$\gamma_i(\hat{x}_i) \sum_{x_{-i} \in X_{-i}} \frac{f_{-i}(x_{-i} \mid x_i)}{b(x_i)} G_i(\hat{x}_i, x_{-i}), \tag{25}$$

We claim that the family $\{f_{-i}(\cdot \mid x_i)/b(x_i) : x_i \in X_i\}$ of vectors satisfies the cone condition for full extraction (Assumption 2 with $f_{-i}(\cdot \mid x_i)/b(x_i)$ taking the role of $f_{-i}(\cdot \mid x_i)$); otherwise, there exist an $x_i \in X_i$ and a nonnegative vector $(\lambda(x'_i))_{x'_i \in X_i \setminus \{x_i\}}$ such that

$$\frac{f_{-i}(x_{-i} \mid x_i)}{b(x_i)} = \sum_{x'_i \in X_i \setminus \{x_i\}} \lambda(x'_i) \frac{f_{-i}(x_{-i} \mid x'_i)}{b(x'_i)}$$

for all $x_{-i} \in X_{-i}$, which implies

$$f_{-i}(\cdot \mid x_i) = \sum_{x'_i \in X_i \setminus \{x_i\}} \frac{\lambda(x'_i) b(x_i)}{b(x'_i)} f_{-i}(\cdot \mid x'_i),$$

contradicting Assumption 2. Now that the cone condition is satisfied, the Farkas lemma implies that there exists function $G_i(\hat{x}_i, \cdot)$ that makes (25) zero if $\hat{x}_i = x_i$ and negative if $\hat{x}_i \neq x_i$. Then the lottery $\gamma_i(\hat{x}_i) g_i(\hat{x}_i, \cdot)$ is obtained via setting

$$g_i(\hat{x}_i, x_i) := \frac{G_i(\hat{x}_i, x_{-i})}{\sum_{s=0}^{\infty} \delta^s \beta(\hat{x}_i, x_{-i}, s)}$$

and picking a scalar $\gamma_i(\hat{x}_i)$ so large that the negative payoff when $\hat{x}_i \neq x_i$ outweighs bidder i's gain from buying the good. Thus, a Crémer-McLean lottery exists for i, as desired. ■

Although the seller can almost fully extract social surplus, search cost makes this task more difficult. In the Crémer-McLean model, for any environment, the size of necessary transfers is given and finite (though the size need not be uniformly bounded when the environment varies). In our proof, by contrast, for any environment where full participation is not socially efficient, the size of transfers needs to be large in the low-probability event that full participation occurs and a bidder's report matches the others' poorly. When the seller reduces the probability of this inefficient event to arbitrarily close to zero, she needs to enlarge the transfers in this event without bound.

5 Discussion

We have studied a single unit auction environment in which the set of bidders is endogenously determined through a dynamic search process. Our main results are that with independent bidders' types, an optimal mechanism amounts to symmetric-information optimal search where the prizes are the virtual utilities. That is, the seller conducts a costly search for the bidder with the highest virtual utility. In traditional optimal auction problems, 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. With correlated bidders' types, the dynamic nature of our model precludes full rent extraction as in the case of a static auction model. Nonetheless, we proved that it is possible to design a mechanism that fully extract the ex ante social surplus with an arbitrarily high probability.

Myerson [9, Ch. 10] has established an elegant virtual utility result in a general model of the traditional mechanism design framework where the set of participants is exogenous. There, the virtual utilities are functions of the dual variables of the designer's constrained optimization problem. Being endogenous, the dual variables usually do not result in explicit characterization of the distortion of asymmetric information. In our model, the designer's problem is a stochastic dynamic programming, and its dual variables may be intractable. Fortunately, the (ex post) virtual utility functions in our model are directly determined by the primitives. That makes the designer's job quite convenient: simply transform the potential bidders' distributions by the virtual utility functions and then plug the transformed distributions into an operation research program that yields a solution for optimal search.

References

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

[2] Jacques Crémer and Richard P. McLean. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257, November 1988.

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

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

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

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

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

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

[9] Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.

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

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

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

[13] Lixin Ye. Optimal auctions with endogenous entry. Mimeo, Department of Economics, Stanford University, 2000.

Index