

Optimal Search Auctions with Correlated Bidder Types*

Jacques Crémer[†] Yossi Spiegel[‡]

September 2005

Abstract

We study the design of profit maximizing single unit auctions under the assumptions that (i) the bidders' values are correlated and (ii) the seller needs to incur search costs when contacting prospective bidders. We show that although in general, the seller cannot fully extract the social surplus as in the case without search costs, it is nonetheless possible to construct a search mechanism that will fully extract the surplus with an arbitrarily high probability.

*This paper grew out of joint research with Charles Z. Zheng.

[†]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>

1 Introduction

It is well known that, in auction environments with risk neutral bidders whose values are correlated, the seller can generically design an optimal auction that implements the socially efficient allocation and gives the seller the entire ex ante social surplus (see Crémer-McLean, 1988). But if the seller must incur search costs in order to contact prospective bidders, then it is in general suboptimal to invite all bidders to participate at once. For instance, if an early bidder turns out to have a sufficiently high value, it may be optimal to sell him the good immediately without incurring further search costs. That is, the seller must design 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 payment that each participating bidder makes. While the sequential nature of the mechanism economizes on the seller's search costs, it may prevent the seller from using the Crémer-McLean lotteries that condition a bidder's payment on all his rivals' reports. Consequently, the seller may be unable to fully extract the surplus as in the case without search costs. The question then is how well can the seller do under these circumstances.

To address this question we consider a private value single-unit auction environment with correlated bidders' types. We assume that in order to contact prospective bidders, the seller needs to incur a cost, to which we refer as *search cost*. We show that although in general the seller cannot fully extract the surplus, it is nonetheless possible to construct a search mechanism that allows the seller to achieve full extraction with a probability arbitrarily close to one. Hence, the sequential nature of the problem, dictated by the presence of search costs, imposes only negligible costs on the seller. Put differently, the seller can do almost as well as if there were no search costs.

Although there is already a small literature that studies optimal search mechanisms, to the best of our knowledge, this paper is the first to deal with correlated bidders' types. Earlier papers have studied the case of independently distributed bidders' types. McAfee and McMillan (1988) characterize profit maximizing search mechanisms when the values of prospective bidders are *i.i.d.* and the seller's cost of contacting bidders is the same across bidders. Burguet (1996) considers 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. Crémer, Spiegel, and Zheng (2005) study the design of profit maximizing search auctions when the bidders are not identical ex ante (in terms of the distribution of their types and the seller's cost of contacting them) and their values are possibly interdependent. Crémer, Spiegel, and Zheng (2003) study a related problem under the assumption that potential bidders do not know their valuations at the outset but can learn them at a cost. A crucial assumption in that paper is that the seller can costlessly contact the bidders before they learn their valuations and offer them contracts.¹

2 The Model

A seller wants to sell an indivisible good to one out of a finite set I of prospective bidders. The seller's value is normalized to zero. Bidder i 's ex post utility from winning the good is equal to the realized value of his type $x_i \in X_i$. A vector of types $x := (x_i)_{i \in I}$ is called a *realized state*. Nature draws a realized state, x , from the set $\times_{i \in I} X_i$ according to a strictly positive joint probability distribution function $f(x)$. Hence, bidders' types are in general correlated. 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 $\delta^{t-s}x_i - \sum_{t=s}^{\infty} \delta^{t-s}p_i^t$ if he gets the good in period $t' \geq s$, and $-\sum_{t=s}^{\infty} \delta^{t-s}p_i^t$ if he never gets the good.

2.1 Search costs²

Initially, bidders do not know the seller's intention to sell the good and are not aware of the auction setting (the rules in the seller's mechanism, who the other bidders are, how their values are distributed, etc.). In order 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*. After being

¹Other papers that study the optimal choice of auction formats with costly information acquisition include Levin and Smith (1994), Ye (2004), Bergemann and Pesendorfer (2001), and Bergemann and Välimäki (2002) (the latter is a general mechanism design problem), and Compte and Jehiel (2000). None of these papers, however, consider search mechanisms as in our paper.

²The material in Sections 2.1-2.3 is based on Sections 2.1 and 2.3-2.5 in Crémer, Spiegel, and Zheng (2005).

contacted by the seller, each bidder i privately learns his type x_i .

The cost c_i has several possible interpretations. First, the good might be very complex; the seller then needs to meet potential bidders in person and describe the good in detail. For example, if the good is the controlling block of a state-owned enterprise, the government might wish to conduct a road show and meet with potential buyers in person. Moreover, the seller may have goals other than profit maximization and he would like to ensure that the bidders meet certain criteria (e.g., ensure that the privatized state-owned enterprise will be controlled by qualified buyer who will be able to manage it properly). Second, although we consider an auction environment, our framework can be easily modified to a procurement environment in which a buyer wishes to buy an indivisible good from a finite set I of potential sellers. 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 interested in outsourcing a custom-made component; in some cases, rather than sending the description to prospective suppliers and asking them for price quotes, it could be more efficient to ask prospective suppliers to describe their manufacturing facilities and then instruct them how to use these facilities in order to produce the good. The supplier can then provide a quote.^{3,4}

2.2 Search mechanisms

When the seller needs to incur search costs in order 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

³For another example, consider a movie producer who is 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 has to offer. Only then can the producer tell which facilities it would need and obtain a price quote.

⁴In 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 (2003) we consider the other extreme case and assume that the buyers need to incur costs in order to learn their types before they can participate in the seller's mechanism. It would be interesting in future research to study the much harder case where both the seller and the buyers need to incur search costs.

plan that, based on the incumbents' messages, specifies whether the seller should stop the mechanism and keep the good or allocate it to one of the incumbent bidders, or whether the seller should continue and invite new bidders (possibly several bidders at once). Coupled with a payment scheme, such a contingent plan is called *search mechanism*.

A search mechanism begins with the seller contacting a set of entrants in period 1. Each entrant privately learns his type and decides whether or not to participate. If an entrant agrees to participate, he signs with the seller a binding contingency contract and sends the seller a message.⁵ Given these messages, either the mechanism stops and the seller keeps the good or allocates it to one of the period-1 entrants, or the mechanism continues to period 2. In the latter case, new entrants are invited, each of them learns his type and decides whether to participate. Each period-2 entrant who decides to participate sends the seller a message. Depending on the messages sent in period 1 and 2, the mechanism either continues and new entrants are invited to participate in period 3, or the mechanism stops and the good is kept by the seller or is allocated to one of the period-1 or period-2 entrants. The mechanism continues in a similar fashion until it stops and the good is allocated. We will allow the possibility of randomization on whether to continue and on the set of new entrants to be invited.

A *revelation search mechanism* is a search mechanism in which each bidder i 's message space is i 's type space. Given a realized state x and assuming that every invited bidder participates and is truthful (which is true in equilibrium), a revelation search mechanism induces the following objects (the formal definitions are straightforward and hence are omitted):

- $E^t(x) :=$ the set of potential bidders who enters the mechanism in period t ;
- $q_i(x) :=$ the probability with which player i consumes the good;
- $p_i(x) :=$ the total payment made by bidder i discounted back to the period at which i enters the mechanism;
- $\tau(x) :=$ the period at which the search terminates.⁶

⁵We assume that the seller cannot force buyers to commit to payment plans before they know their types, otherwise, prospective buyers would be besieged by dishonest sellers selling them useless goods. This idea can be formalized by adding a preliminary stage to the model in which nature selects the seller's type, which can either be bona fide (as we assume here) or fake.

Hence a revelation search mechanism can be denoted by $((E^t)_{t=1}^\infty, (p_i, q_i)_{i \in I})$. The sequential nature of a search mechanism imposes the following constraints on the above objects:

1. E^1 is constant on $X := \times_{i \in I} X_i$, i.e., the set of entrants in period 1 is determined before any message is sent.
2. If two realized states x and x' generate the same set of entrants up to period $t \geq 1$, i.e., $E^s(x) = E^s(x')$ for all $s \leq t$, and if $x_i = x'_i$ for all incumbents $i \in \cup_{s=1}^t E^s(x)$, then:
 - a. The set of entrants in period $t + 1$ is the same for x and x' and is a subset of potential bidders who have not yet entered, i.e., $E^{t+1}(x) = E^{t+1}(x') \subseteq I \setminus \cup_{s=1}^t E^s(x)$.
 - b. If the mechanism stops at period t , then the probability with which bidders win the good and how much they pay is the same for x and x' , i.e., $q_i(x) = q_i(x')$ and $p_i(x) = p_i(x')$ for all potential bidder i .
3. Bidders who do not participate in the mechanism, do not have a chance of winning the good and do not pay anything, i.e., $q_i(x) = p_i(x) = 0$ for all $i \notin \cup_{t=1}^\infty E^t(x)$.

A *search procedure*, denoted $((E^t)_{t=1}^\infty, (q_i)_{i \in I})$, 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 the bidders' payments.

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 following result will allow us to restrict attention to revelation search mechanisms. The proof of the lemma is similar to the proof of the standard revelation principle and hence is omitted.

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

2.3 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 extracts the entire surplus, then the seller's expected profit discounted to period 1 is equal to

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

Note that $q_i(x) (x_i - x_0)$ can be summed over all potential bidders (even those who did not enter the mechanism before it stopped) because by constraint 3 in the previous subsection, $q_i(x) = 0$ for all nonparticipating bidders.

In traditional search theory there is no *asymmetric* information once the search cost has been incurred. Hence, optimal search amounts to maximizing $\Pi((E^t)_{t=1}^\infty, (q_i)_{i \in I})$ over all search procedures. 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. Consequently, the seller must design a search mechanism that induces each bidder to reveal his type truthfully. We will say that a search procedure $((E^t)_{t=1}^\infty, (q_i)_{i \in I})$ is *optimal* if $\Pi((E^t)_{t=1}^\infty, (q_i)_{i \in I})$ is maximized over all search procedures subject to the bidders' incentive feasibility constraints.

3 The impossibility of full extraction of the surplus

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, 1988). The presence of search costs, however, renders such full extraction difficult because a search procedure continues or stops depending on the incumbents' reports. Hence an incumbent bidder can prevent the entry of rivals and therefore make it impossible to use the appropriate lotteries.

The following example illustrates the source of the difficulty. Suppose that there are two ex ante identical bidders, 1 and 2, whose types, x_1 and x_2 , are drawn from the set

$\{L, M, H\}$, where $L < M < H$. The discount factor is one, and the search cost of contacting each bidder is $c > 0$. The joint probability distribution 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; 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 that

$$f_{LL} + f_{MM} + f_{HH} > c, \quad (2)$$

where $f_x := f_{xL} + f_{xM} + f_{xH}$ for $x = L, M, H$. Note that since the discount factor is one, there is no loss of efficiency to invite only one of the 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

$$\frac{f_{LM}M + f_{LH}H - (f_{LM} + f_{LH})L}{f_L} < c < \frac{f_{MH}(H - M)}{f_M}, \quad (3)$$

it is optimal to stop if $x_1 = L$ but continue if $x_1 = M$.⁷ Therefore, whenever (3) holds, the mechanism stops in period 1 if $x_1 \in \{L, H\}$ and bidder 2 is never contacted. Consequently, when $x_1 \in \{L, H\}$, bidder 1's payment cannot depend on bidder 2's report and by incentive compatibility must be the same. Since interim participation constraint when $x_1 = L$ implies that bidder 1's payment cannot exceed L , it is clear that the seller cannot fully extract the social surplus when $x_1 = H$.

From the example it is immediate that the following result holds:

Proposition 1 *If the socially efficient search procedure invites a single bidder i in period 1, and if there are at least two possible realized values of bidder i 's type at which the efficient*

⁷If $x_1 = L$, stopping yields a social surplus L while contacting bidder 2 and allocating the good to the higher-value bidder yields expected 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 expected surplus of $\frac{(f_{LM} + f_{MM})M + f_{MH}H}{f_M} - c$. When (3) holds, then $L > \frac{f_{LL}L + f_{LM}M + f_{LH}H}{f_L} - c$ and $M < \frac{(f_{LM} + f_{MM})M + f_{MH}H}{f_M} - c$.

procedure calls for awarding the good to bidder i without further search, then the seller cannot fully extract the maximum social surplus in the symmetric-information search problem.

4 Almost full extraction of the surplus

Although full extraction of the surplus 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 continues with a positive probability, thereby eliminating the first entrant's ability to exclude rivals.

Let us start with the previous example, assuming that both (2) and (3) are satisfied and assuming that the above joint probability matrix satisfies the cone condition for full extraction (Crémer and McLean, 1988; stated below as Assumption 1). 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:

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

For instance, $\Pr(x_1 = L \mid x_2 = H) = \frac{\epsilon f_{LH}}{\epsilon f_{LH} + f_{MH} + \epsilon f_{HH}}$. 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. Since the seller fully extracts the social surplus when the efficient procedure is implemented, he also obtains the maximum social surplus with probability of at least $1 - \epsilon$.

Recall that $f(x)$ is the joint prior probability of the realized state x , with $f(x) > 0$ for all realized states x . Let $X_{-i} := \times_{j \neq i} X_j$. For any $x_{-i} \in X_{-i}$, let $f_{-i}(x_{-i} | x_i)$ 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, $f_{-i}(x_{-i} | x_i)$ is well defined. Note that $f_{-i}(\cdot | x_i)$ 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 (1988).

Assumption 1 *For any bidder i and any $x_i \in X_i$, the vector $f_{-i}(\cdot | x_i)$ does not belong to the cone generated by the vectors in the family $\{f_{-i}(\cdot | x'_i) : x'_i \in X_i \setminus \{x_i\}\}$, i.e., there does not exist a family $\{\rho_i(x'_i; x_i)\}_{x'_i \in X_i \setminus \{x_i\}}$ such that*

$$f_i(x_{-i} | x_i) = \sum_{x'_i \in X_i \setminus \{x_i\}} \rho_i(x'_i; x_i) f(x_{-i} | x'_i) \text{ for all } x_{-i} \in X_{-i}.$$

As in the above 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 every entrant always assigns a positive posterior probability to any possible realized state.

Theorem 1 *Given Assumption 1, 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 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. 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 $t = 1, 2, \dots$. Given i 's report, \hat{x}_i , suppose that i 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.⁸ Derived from the design of the mechanism, $a(x_{-i})$ is commonly known. Likewise, given any x_{-i} and $s = 0, 1, 2, \dots$, 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

$$b(x_i) := \sum_{x'_{-i} \in X_{-i}} a(x'_{-i}) f_{-i}(x'_{-i} | x_i);$$

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

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

$$\gamma_i(\hat{x}_i) \sum_{x_{-i} \in X_{-i}} \frac{f_{-i}(x_{-i} | x_i)}{b(x_i)} G_i(\hat{x}_i, x_{-i}), \quad (4)$$

We claim that the family $\left\{ \frac{f_{-i}(\cdot | x_i)}{b(x_i)} : x_i \in X_i \right\}$ of vectors satisfies the cone condition for full extraction (Assumption 1 with $\frac{f_{-i}(\cdot | x_i)}{b(x_i)}$ taking the role of $f_{-i}(\cdot | x_i)$); 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

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

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

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

thereby contradicting Assumption 1. Since the cone condition is satisfied, the Farkas lemma implies that there exists a function $G_i(\hat{x}_i, \cdot)$ that makes (4) 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. ■

⁸ Actually $a(x_{-i})$ depends only on the reports of the incumbents before i enters.

5 Discussion

We have studied a single-unit auction environment with correlated bidders' types in which the set of bidders is endogenously determined through a dynamic search process. Our main results are that (i) the dynamic nature of our model precludes full extraction of the surplus as in the static auction case, and (ii) despite the difficulties that arise due to the seller's need to incur search costs when contacting prospective bidders, it is nonetheless possible to design a mechanism that fully extracts the ex ante social surplus with an arbitrarily high probability.

It should be noted however that although the seller can almost fully extract the social surplus, the presence of search costs makes this task considerably harder. The reason for this is as follows. In the Crémer-McLean model, the size of necessary transfers is given and finite for any environment (albeit it need not be uniformly bounded when the environment varies). In the proof of Theorem 1, 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, he needs to increase the transfers in this event without a bound.

6 References

Bergemann, Dirk and Martin Pesendorfer (2001), "Information Structure in Optimal Auctions," Cowles Foundation Discussion Paper No. 1323.

Bergemann, Dirk and Juuso Välimäki (2002), "Information Acquisition and Efficient Mechanism Design," *Econometrica*, **70**, 100- 1033.

Burguet, Roberto (1996), "Optimal Repeated Purchases When Sellers are Learning about Costs," *Journal of Economic Theory*, **68**, 440-455.

Compte, Olivier and Philippe Jehiel (2000), "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

Crémer J. and R. McLean (1985), ”Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent,” *Econometrica*, **53**, 345-361.

Crémer J. and R. McLean (1988), ”Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, **56**, 1247-1257.

Crémer, Jacques, Yossi Spiegel, and Charles Zheng (2003), “Optimal Selling Mechanisms with Costly Information Acquisition,” Working paper, Institute d’Economie Industrielle, Toulouse.

Crémer, Jacques, Yossi Spiegel, and Charles Zheng (2005), “Optimal Search Auctions,” Mimeo.

Levin, Dan and James L. Smith (1994), “Equilibrium in Auctions with Entry,” *American Economic Review*, **84**, 585-599.

McAfee, Preston and John McMillan (1988), “Search Mechanisms,” *Journal of Economic Theory*, **44**, 99-123.

Ye, Lixin (2004), “Optimal Auctions with Endogenous Entry,” *Contributions to Theoretical Economics*, **4**, Article 8.