Search with Adverse Selection*

Stephan Lauermann Asher Wolinsky

This version: July 2008

Abstract

This paper explores the "lemons problem" with search. A buyer samples sellers to procure a service. The sellers' cost of providing the service depend on the type of the buyer. We discuss the implications of changes in the size of the frictions (search costs) and the implications of giving sellers informative signals about the type of the buyer. As we show, frictions have a positive role in this market as they shield sellers from the winner's curse. We demonstrate the strength of the winner's curse in a search market by showing that even if sellers receive arbitrarily informative signals, equilibria with complete pooling exist. We compare our results for a search market with results for auctions, showing that in a search market stronger conditions are necessary to ensure information revelation.

The bargaining process is a critical part of the model. Due to the information asymmetry, we cannot use the simple surplus sharing solutions that are common in the search literature. Instead, we assume that the buyer has all the bargaining power and offers a general *mechanism* which the seller can either accept or reject. Thus, we model bargaining as a principal-agent problem, with the buyer being the *informed principal* as in Myerson (1983).

Very preliminary and Incomplete. Please do not post publicly.

JEL Classifications: D44, D82, D83

Keywords: Adverse Selection, Winner's Curse, Search Theory, Auctions, Information Aggregation.

*The title is preliminary and subject to change. The full paper is yet to be written and might differ substantially from the current version. Acknowledgements to be added later.

1 Introduction

This paper looks at model of search with adverse selection. An agent that we call the buyer samples sequentially alternative trading partners, sellers, for a transaction that involves information asymmetry. To have a story in mind, one can think of the buyer as someone who is applying for credit and who is seeking offers from potential lenders. The existence of good and bad creditors turns this into a setting with adverse selection. The main objective of this paper is to understand how the combination of search activity and information asymmetry affects prices and welfare.

If sellers cannot tell good and bad buyers apart, equilibria will involve pooling. But what happens if sellers receive some noisy information (a signal) about the type of the buyer? We want to understand under which conditions such signals will lead to a separation of the buyer's types.

If a buyer faces high search cost, perfect separation is not possible. By chance the good buyer might sample a seller who mistakes him for a bad type. Similarly, a bad buyer might happen to sample a seller who mistakes him to have a good type. Indeed, a buyer with a bad type might deliberately try to sample many seller to find some unlucky one who mistakes him for a good type, allowing the buyer to get a good deal. In equilibrium, sellers will take this behavior into account when interpreting their own information. Thus, the precision of the seller's information is *endogeneous* in our model.

Our main result concerns a situation in which search cost become small, i.e., we are looking at a limit. We say the limit involves complete pooling if good and bad buyers trade at the same prices. The limit is perfectly revealing if buyers get the same price they would get with perfect information. We show the limit is perfectly revealing if and only if there are arbitrarily informative signals[1] *and* the tail of the distribution of the signals is sufficiently thick. If arbitrarily informative signals do not exist, the limit necessarily involves complete pooling. This result demonstrates that separation is hard to achieve in a model with search. The reason for this negative result is excessive search of the bad types, diminishing the value of information and excerbating the winner's curse for the seller.

We compare our result to a setting in which a buyer can commit to a procurement auction and we look at the case in which the number of bidders become large. In a procurement auction the limiting outcome never involves total pooling. Furthermore, if arbitrarily informative signals exist, the limiting outcome will be perfectly informative as shown by Milgrom (1979) and Wilson (1977).[2]

[1] Signals exist that are so informative that they are arbitrarily close to reveale the state.

[2] Duggan and Martinelli (2001) find that the existence of arbitrarily informative signals playes a similar role in a model of voting in juries when the size of the jury increases.

In contrast, in a model with search, the outcome can involve complete pooling even with arbitrarily informative signals.

(Our second set of results, not reported in this draft, concerns situations with small but not vanishing search cost. We point out that more substantial search frictions may result in lower prices and higher welfare. This is in contrast to common intuition emanating from other search models that lower search frictions are associated with lower prices and higher welfare. The source of the difference is information asymmetry of the sort present in this model. Lower search costs result in relatively more search by the bad types hence strengthening the winner's curse. This effect is absent of course from search models that do not incorporate this information asymmetry and where search often involves a positive externality and hence lower frictions tend to enhance welfare through both a direct and an indirect effect. Thus, not only can less information be good for efficiency, higher frictions can be good as well.)

Switching the roles of buyers and sellers, one can interpret our model as a "lemons market" in which a sellers of a used car is searching for a buyer. Another interpretation is that of a labor market with adverse selection, in which an unemployed worker of unkown quality searches for a job. Since unemployed worker of different types generate different signals, the search duration and the final wage will be correlated in a systematic manner. While not reported here, it will be the case in many equilibria that the bad type searches longer than the good type (because it is harder to generate favorable signals) but still ends up getting a lower price/wage.

Our model is this: A buyer searches sequentially among sellers to obtain a service. The value of the service to the buyer is commonly known. The buyer incurs a cost $s > 0$ ("search cost") to sample a seller. A seller's cost of providing the service, c_w, is the same for all sellers and it depends on an underlying state $w \in \{L, H\}$ with $c_H > c_L$. The state w is known to the buyer but not to the sellers. We shall call w also the *type* of the buyer.

At the beginning of every sampling round, the buyer draws one seller at a cost s. The seller receives a signal that is correlated with the state. The signal is jointly observed by the buyer and the seller. Then, the buyer and the seller bargain over the terms of trade, to be described below. If they reach an agreement and trade, the game is over. If they do not trade and if the buyer chooses to proceed, the next round starts according to the same rule and the buyer samples another seller at cost s.

The bargaining process that takes place after a seller is sampled by the buyer is a critical part of the model. Due to the information asymmetry, we cannot use the simple surplus sharing solutions that are common in the search literature with symmetric information. A simple surplus sharing rule is characterized by a number $\beta \in [0, 1]$ such that the buyer receives a share β of the surplus. With complete information, a surplus sharing rule is equivalent to a game in which, with probability β, the buyer has all bargaining power and makes a take-it-or-leave-it price offer to the seller (and with probability $(1-\beta)$ the seller makes such an offer).

We extend this simple game to a setting with asymmetric information and interdependent valuations. We assume that the buyer has all the bargaining power. In our setting, however, offering a simple price is no longer an optimal mechanism. Instead, the buyer offers a *mechanism* which the seller can either accept or reject. Thus, we model bargaining as a principal-agent problem, with the buyer being the *informed principal* as in Myerson (1983), proposing a trading mechanism to the seller (agent). Since the principal has information that affects the preferences of the agent, the mechanism proposal game is a signaling game. In general, the game suffers from large multiplicity of equilibria due to the freedom of specifying beliefs off the path. We therefore employ a number of refinements. Given these refinements, we characterize the set of equilibrium mechanisms. The identified mechanisms are interim efficient and the full surplus of trade is extracted by the buyers. By allocating bargaining power randomly and allowing a seller to be the proposer of a mechanism with some probability as well, we could capture situations with intermediate degrees of bargaining power as well.

Our main result concerns the limit of the equilibrium outcomes when s becomes small. Let F_w denote the distribution of the sellers beliefs in state w, conditional on their signal. We show that in the limit of every equilibrium the two types of buyers will trade at a price equal to the true costs c_w if and only if the appropriately defined tail of F_w is thick enough.

In the next section we describe the model. Then we characterize the set of equilibrium mechanisms. We show that an equilibrium is offered as part of some equilibrium only if it maximizes the expected payoff of the low cost buyer, subject to the constraint that the mechanism must be individually rational for the seller and the high cost buyer (i.e., the high cost buyer prefers to offer the same mechanism rather than any other mechanism, given the beliefs of the seller.) The provision and analysis of the informed principal under adverse selection is a significant contribution of the paper and should be fruitfully applicable to similar models. Existing models usually circumvent the difficulties inherent in the analysis of the signaling game by asssuming that only the uninformed party makes the offer.

In our model, giving bargaining power to the seller would lead to monopolistic prices, i.e., prices would be equal to the value of the service. However, since there would be no search on the equilibrium path, the outcome could in principle be more efficient. Therefore, in our model, giving bargaining power to the *uninformed* party can be good for efficiency while giving bargaining power to the informed party can be bad.

2 The Model

A buyer searches sequentially among sellers to obtain a service. To have a story in mind, one may think of a procurement scenario in which the buyer is seeking to fix a problem (repair or cure) and samples service providers sequentially to obtain bids[3]. The value of the service is commonly known and denoted by u. The buyer incurs a cost $s > 0$ ("search cost") to sample a seller. A seller's cost of providing the service, c_w, is the same for all sellers and it depends on an underlying state $w \in \{L, H\}$ with $c_H > c_L$. The prior probabilities of L and H are g_L and g_H respectively. The state w is known to the buyer but not to the sellers. We shall call w also the *type* of the buyer. The value of the service u is sufficiently larger than $c_H + s$ (the cost of the service and the cost of finding a seller) so that both types of the buyer would like to participate.

At the beginning of every sampling round, the buyer draws one seller at a cost s. The seller receives a signal $x \in X \subset [0, 1]$ that is correlated with the state. The distribution of x given w is F_w and we allow for atoms. (In particular, in a leading example X has only two numbers as its elements, $X = \{l, h\}$.) F_w satisfies the montone likelihood ratio property and a low signal is indicative of the low state. The buyer observes the signal of the seller. Then, the buyer offers a direct mechanism M to the seller. The seller can either accept or reject the mechanism. If the mechanism is accepted, the buyer reports his type and the mechanism implements the prescribed allocation. If the mechanism is accepted and if trade happens, the game stops. If either the mechanism is rejected or if the mechanism prescribes no trade, the buyer can choose to stop the game.[4] If the buyer chooses to proceed, the next round starts according to the same rule and the buyer samples another seller at cost s.

If the buyer transacts at a price p after having sampled n sellers, his payoff is $u - p - ns$. The payoff of the seller who agreed to the transaction is $p - c_w$. The payoff of all other sellers is zero. The realized surplus is $u - c_w - ns$.

A collection of strategies - the mechanism offer M, acceptance decision A, and reporting decision R - and beliefs β of the seller is called a constellation σ. A direct mechanism M is a vector $[p_L, q_L, p_H, q_H]$, where p_w, q_w are the trading price and the trading probability conditional on a report $R = w$.[5] $M(x, w)$ describes the mechanism offered by a buyer w if the seller has a signal x and a

[3] Alternatively, one may reverse the roles of what we call buyer and sellers to obtain an even more standard story of sale of an object of uncertain quality w.

[4] A seller always accepts a price above c_H. Since $u > c_H + s$, this implies that it is always worthwhile to continue and the buyer will never stop sampling. We therefore do not include a stopping decision in the formal analysis.

[5] The price is paid conditional on trading, i.e., the expected transfer given a mechanism M and a report R is $t_R = q_R p_R$. This is without loss of generality relative to specifying transfers if the trading probability is positive whenever transfers are nonzero. This will the case in equilibrium.

mechanism offer strategy is $M(\cdot,\cdot):\{L,H\}\times X\rightarrow\mathbb{R}_{+}^{4}$.[6] Given a mechanism offer M, the seller believes that the probability of the high state is $\beta(M,x)$, so beliefs are $\beta(\cdot,\cdot):X\times\mathbb{R}_{+}^{4}\rightarrow[0,1]$. $A(M,x)$ describes the acceptance decision by a seller of type x, $A(\cdot,\cdot):\mathbb{R}_{+}^{4}\times X\rightarrow[0,1]$, where A is the acceptance probability. If the mechanism is accepted, $R(w,x,M)$ describes the report conditional on the state w, the signal x, and the offered mechanism M, $R:\{L,H\}\times X\times\mathbb{R}_{+}^{4}\rightarrow\{L,H\}$. We will define an equilibrium as a constellation in which the strategies are mutually *optimal* and the beliefs are *consistent*. By the inscrutability principle, there is no loss of generality in assuming that both buyers offer the same mechanism, the mechanism is accepted, and reports are truthful. Let us define these requirement precisely.

Given a constellation σ describing the behavior of the other players and their beliefs, expected payoffs of the buyer who samples a seller with signal x and who uses strategy M,R is recursively defined. The payoff is the probability of trading with the current seller times the expected profit conditional on trading plus the expected continuation payoff if no trade happens minus the search costs:

$$\begin{aligned}U_w(M,R,x,\sigma) &= A(M(x),x)\,q_{R(w,x,M(x))}^{M(x)}\left(u-p_{R(w,x,M(x))}^{M(x)}\right)\\ &\quad+\left(1-A(M(x),x)\,q_{R(w,x,M(x))}^{M(x)}\right)\int_x U_w(M,R,x,\sigma)-s\end{aligned}$$

where q_R^M is the trading probability in mechanism M given report R and similarly for p_R^M. Let $V_w(\sigma)$ be the expected payoff of the buyer who uses the strategies prescribed by σ. The payoff of a seller with signal x who received an offer M is equal to the probability to accept the contract times the expected profit from the contract conditional on the high cost and the low cost buyer, respectively, weighted by the relative probabilities

$$\begin{aligned}\pi(M,A,x,\sigma) &= A(M,x)\left(\beta(M,x)\,q_{R(H,x,M)}^{M}\left(p_{R(H,x,M)}^{M}-c_H\right)\right.\\ &\quad\left.+(1-\beta(M,x))\,q_{R(L,x,M)}^{M}\left(p_{R(L,x,M)}^{M}-c_L\right)\right).\end{aligned}$$

Let $\beta_0(x,\sigma)$ denote the prior belief of a seller with signal x who is sampled by a buyer. The probability of being sampled might depend on the state, since the two types of buyer might sample a different number of seller in expectation. Let $\theta(\sigma)$ denote the ratio of the number of sellers who are sampled in expectation,

$$\theta(\sigma)=\frac{E[\#\text{ of sellers sampled}|w=L,\sigma]}{E[\#\text{ of sellers sampled}|w=H,\sigma]}.$$

If the low cost buyer L samples many more sellers than the high cost buyer, a seller who is sampled should update towards the low state. As shown in the

[6] We disregard all measurability issues throughout the paper,e.g., we do not restrict the set of mechanisms by requiring $M(\cdot)$ to be measureable.

appendix, the prior of the seller can be defined as

$$\beta_0(x,\sigma) = \frac{g_H dF_H(x)}{g_H dF_H(x) + g_L dF_L(x)\,\theta(\sigma)}.$$

In equilibrium, both buyers will offer the same mechanism and thus, the beliefs of a seller following an on-equilibrium mechanism offer are just $\beta(x,M) = \beta_0(x,\sigma)$. Off-equilibrium beliefs are not restricted.

The acceptance decision by the seller is sequentially rational if he plans to accept mechanisms that lead to strictly positive profits and if he plans to reject mechanisms that lead to negative profits, given his beliefs $\beta(M,x)$, i.e.,

$$A^*(M,x) = \begin{cases} 1 & \text{if} \quad \pi(M, A=1, x, \sigma^*) > 0 \\ 0 & \text{if} \quad \pi(M, A=1, x, \sigma^*) < 0 \end{cases}.$$

By the inscrutability principle (Myerson, 1983), we can restrict attention to equilibria in which both types of the buyer offer the same, direct mechanism, the mechanism is accepted, and reports are truthful. Every equilibrium outcome of a larger game in which the buyer can offer more complex mechanism is equivalent to an outcome of an equilibrium in which both types of buyers offer the same, direct mechanism that is incentive compatible and individually rational. We will therefore drop the dependency of the offer strategy $M(w,x)$ on w during the analysis.

Note that we define strategies to be history independent, i.e., the buyer can condition his mechanism offer only on his own type and the signal of the seller and the reporting strategy may depend in addition on the offered mechanism. The seller's acceptance strategy depends only on the signal and on the mechanism offer (since sellers do not observe anything else). Our basic equilibrium definition is therefore essentially that of a Markov Perfect equilibrium:

Definition 1 *A constellation σ^* is an inscrutable equilibrium if*

1. $M^*(x,w)$ and R^* are *optimal,* $M^*, R^* \in \arg\max U_w(M,R,x,\sigma)$.
2. β^* is derived from *Bayes Rule* whenever applicable.
3. A^* is *sequentially rational.*
4. M^* is *accepted* and *reporting is truthful,* $A^*(M^*,x) = 1$ and $R^*(w,x,M^*) = w$.
5. The equilibrium is *inscrutable,* $M^*(x,H) = M^*(x,L)$.

As indicate before, we impose refinements. We discuss the implications of these refinements in the next sections.

A mechanism M is said to be *undominated* for x given σ if, conditional on continuation payoffs $V_w(\sigma)$ and a seller x with belief $\beta_0(x,\sigma)$, there is no other mechanism M' that would make both types of the buyer and the seller (in expectation) strictly better off, i.e., if there is no M' such that, given acceptance and truthful reports, the payoff would be higher for both types of buyers,

$$q_w^{M'}\left(u-p_w^{M'}\right)+\left(1-q_w^{M'}\right)V_w > q_w^M\left(u-p_w^M\right)+\left(1-q_w^M\right)V_w, \quad w\in\{L,H\}$$

and, given acceptance and truthful reports by both buyers, the payoff of the seller would be higher

$$\beta_0 q_H^{M'}\left(p_H^{M'}-c_H\right)+(1-\beta_0)\,q_L^{M'}\left(p_L^M-c_L\right) > \beta_0 q_H^M\left(p_H^M-c_H\right)+(1-\beta_0)\,q_L^M\left(p_L^M-c_L\right).$$

Beliefs following off equilibrium mechanism offers are said to satisfy "*Divinity*" if they put more probability on the type of buyer who has a stronger incentive to offer the mechanism.[7] We need to define a "stronger incentive." Suppose continuation payoffs are V_w and suppose the buyer believes a mechanism M' is accepted with probability α, then his expected payoff when reporting R is

$$\alpha q_R^{M'}\left(u-p_R^{M'}\right)+\left(1-\alpha q_R^{M'}\right)V_w-s.$$

Let $\alpha_w(M',x,\sigma)$ be the minimal, strictly positive probability at which the payoffs from M' are higher than the payoff from following the strategies prescribed by σ, given an optimal reporting strategy

$$\alpha_w=\inf\left\{\alpha\in(0,1],\infty|\alpha q_R^{M'}\left(u-p_R^{M'}\right)+\left(1-\alpha q_R^{M'}\right)V_w-s>U_w^*\right\},$$

with $U_w^*=U_w\left(M^*(x),R^*,x,\sigma^*\right)$. The definition is such that $\alpha_w=\infty$ if there is no $\alpha>0$ for which acceptance of M' makes w strictly better off.

We say that type w has stronger incentives to deviate from M^* to M' if the cutoff $\alpha_w(M',x,\sigma)$ is lower for w than for the other type. If this is true, then, whenever the deviation is profitable for the other type given some believe about the acceptance probability, it is certainly profitable for w. "Divinity" requires that in this case the posterior of the seller should put no less probability on w than his prior. Beliefs $\beta(x,M')$ satisfy divinity given σ if

$$\begin{array}{rcl}\beta(x,M') & \geq & \beta_0(x,\sigma)\ \text{ if } \alpha_H(M',x,\sigma)<\alpha_L(M',x,\sigma)\\ \beta(x,M') & \leq & \beta_0(x,\sigma)\ \text{ if } \alpha_H(M',x,\sigma)>\alpha_L(M',x,\sigma)\end{array}$$

[7] Since at this point it is *not* a direct application of the original definition but an application of its monotonicity notion, Divinity is put in quotation marks. To apply Divinity directly, we will have to define a reduced, static game between a given buyer and seller, using the continuation payoffs to define static payoffs in a way that preserves the strategic elements of the original (dynamic) game.

Divinity (rather than stronger refinements like D1/D2) is used because it makes the construction of equilibrium easier, for example assigning the belief $\beta_0(x, \sigma)$ off the equilibrium path would ensure that an equilibrium satisfies Divinity.

Divinity (as well as most of the other refinements) for signaling games relies on a single crossing condition on preferences. The condition does hold in our setup if the expected payoff of the low cost buyer is higher than the expected payoff of the high cost buyer. We restrict attention to equilibria in which the payoffs $V_H(\sigma)$ and $V_L(\sigma)$ are ordered in this way. Thus, we rule out a class of pooling equilibria in which both types of buyers trade at the same price.[8]

Here is the equilibrium definition which we will use most of the time. Whenever we use the term equilibrium without qualification, we mean an undominated, monotone equiilibrium:

Definition 2 A *constellation* σ^* *is an undominated, montone equilibrium if* σ^* *is an unscrutable equilibrium and if*

1. $M^*(x)$ is *undominated* given σ for all x.
2. $\beta^*(x, M)$ satisfies *Divinity given* σ for all x and M.
3. Payoffs are *monotone*, $V_L(\sigma^*) > V_H(\sigma^*)$.

3 Existence and Preliminary Observations

In this section we discuss and show existence of an undominated, monotone equilibrium. Let us characterize the set of mechanisms $M(x)$ that satisfy the equilibrium definition for given continuation payoffs $V_w(\sigma)$ and prior beliefs $\beta_0(x, \sigma)$.

Given beliefs β_0, the expected cost of a seller with signal x is

$$E_0[c|x, \sigma] = \beta_0 c_H + (1 - \beta_0) c_L.$$

A subscript zero refers to the evaluation of the expectation at the "prior," accounting for the information contained in the signal x and being sampled, but *not* accounting for the information contained in the mechanism offer.

We say a mechanism M is *feasible* relative to continuation payoffs V_w and beliefs β_0 if the mechanism is individually rational for the buyer (better than no trade), truthful reporting is incentive compatible, and if it is individually rational for the seller to accept the mechanism, given truthful reports. Suppose

[8] Intuitively, this restriction makes it harder to find pooling equilibria and thus strengthens the result that separation is unlikely.

a buyer w has sampled a seller with signal x and suppose the continuation payoff of the buyer is V_w. Then reporting truthfully is optimal if

$$q_w^M\left(u-p_w^M\right)+\left(1-q_w^M\right)V_w \geq q_{w'}^M\left(u-p_{w'}^M\right)+\left(1-q_{w'}^M\right)V_w, \quad \forall w' \in \{L,H\}.$$

Acceptance of a mechanism is profitable if expected profits are positive

$$\beta_0 q_H^M\left(p_H^M-c_H\right)+\left(1-\beta_0\right)q_L^M\left(p_L^M-c_L\right) \geq 0.$$

Given continuation payoffs $V_w(\sigma)$ and beliefs $\beta_0(x,\sigma)$, let $\hat{M}(x,\sigma)$ be the set of feasible mechanisms that maximize the expected payoff of the low cost buyer, subject to the constraint that the high cost buyer has an incentive to offer the mechanism as well:

$$\begin{aligned}
\hat{M}(x,\sigma) \quad = \quad & \arg\max_{[p_L,q_L,p_H,q_H]}\left(q_L\left(u-p_L\right)+\left(1-q_L\right)V_L(\sigma)\right) \\
& s.t.\ \text{M is feasible given } V_w(\sigma) \text{ and } \beta_0(x,\sigma) \\
& \qquad q_H\left(u-p_H\right)+\left(1-q_H\right)V_H(\sigma) \geq u-c_H \quad \text{(IR H-Buyer)}
\end{aligned}$$

The last inequality implies that the H buyer does not want to deviate from M to a mechanism that prescribes trade at c_H for sure. We will argue that this is sufficient to ensure that H does not find any other mechanism profitable, given suitably chosen beliefs by the seller.

As shown in the appendix, a mechanism M can be part of an equilibrium only if it maximizes the payoff of the L buyer as defined before:

Theorem 1 *A mechanism M is part of an equilibrium σ^* only if $M \in \hat{M}(x,\sigma^*)$.*

The intuition behind the theorem is as follows. First, sellers must not receive positive profits. Otherwise, the low cost buyer could use the slack in the seller's IR constraint to make himself better off as shown in the proof. Second, suppose the theorem would not hold and there would be some feasible mechanism M' such that the low cost buyer prefers the allocation (p'_L, q'_L) strictly to the allocation in the original mechanism (p_L, q_L). If the high cost buyer strictly prefers the allocation from M' to the one from M as well, then moving from M to M' would make both types of buyers strictly better off. Increasing the prices slightly would not change this, but would strictly increase the profit of the seller. Thus, M' would *dominate* M and hence M cannot be part of an equilibrium. If, on the other hand, the high cost buyer prefers the original mechanism M (weakly) over M', then, following the deviation to M', Divinity requires that the seller updates towards the low cost buyer. Again, increasing prices slightly (in particular, increasing p_L) would imply that the seller's expected profits from accepting M' are strictly postive and thus, the deviating mechanism M' would be accepted. Thus, M' would be a profitable deviation for the low cost buyer.

The next lemma uses the theorem to characterize equilibria. Define Δ to be the surplus for the H buyer, $\Delta = (u - c_H) - V_H$. Let x^* be the signal such that the available surplus in expectation is equal to the continuation payoff of the good buyer, $x^* = u - E[c|x^*] - \max\{(1-\beta)\Delta, 0\} = V_L$. Also, define $q_L^C = \frac{(u - c_H - V_H)}{(u - c_L - V_H)}$. Then

Lemma 1 *In equilibrium, for all $x < x^*$, $M(x) = [1, E[c|x], 1, E[c|x]]$, $E[c|x] = E_0[c|x, \sigma]$. Furthermore,*
a) if $\Delta > 0$, and $x > x^*$ then $M(x) = [q_L^C, c_L, 1, c_H]$
b) if $\Delta < 0$, and $x > x^*$ then $M(x) = [0, c_L, 0, c_H]$
c) if $\Delta = 0$, and $x > x^*$ then $M(x) = [0, c_L, q_H, c_H]$ for any $q_H \in [0, 1]$.

We do not characterize mechanisms at the point x^*. But note that the behavior at x^* does not affect equilibrium if the signals distribution does not contain atoms.

The Lemma allows to characterize equilibrium mechanism by three numbers, the cutoff x^*, and the trading probabilities q_H and q_L^C (beliefs $\beta(x^*)$ and the expected cost $E[c|x]$ can be inferred). We conjecture that a simple fixed point theorem on these numbers proves existence of equilibrium.

4 Analysis for small s

In the first subsections we look at an example with two signals. First, we show that with small s the equilibrium in the search market involves complete pooling. Second, we compare this outcome to that of an auction, in which the buyer can commit to sample bids from a given number N of sellers. We will see that even with a large number of bidders, complete pooling is not an equilibrium. Third, we discuss the implications of the refinements. For each refinement, we show an example of an equilibrium and discuss how this equilibrium violates our results.

4.1 Discrete Signals

Suppose there are only two signals, $X = \{l, h\}$, with $l < h$. The low signal is indicative of the low state. The probabilities that a seller receives a signal l are $f_{l|L}$ and $f_{l|H}$, with $1 \geq f_{l|L} > f_{l|H} \geq 0$.

Take a sequence of search costs $\{s_k\}$, $s_k \to 0$. For each s_k, pick some equilibrium σ_k^*. Let $E_w[p|\sigma_k^*]$ be the expected price paid by the L buyer in expectation in state w, given s_k and σ_k^*. Let $\bar{p}_w$ be limit price, $\bar{p}_w = \lim_{s_k \to 0} E_w[p|\sigma_k^*]$ (if it exists). (The appendix contains a sketch of the proof of the result.)

Theorem 2 *The limit is separating, $\bar{p}_L \neq \bar{p}_H$ if and only if the low signal is revealing, $f_{l|H} = 0$. If $f_{l|H} > 0$, then the limit involves complete pooling at the ex ante expected price, i.e.,*

$$\bar{p}_L = \bar{p}_H = g_H c_H + g_L c_L.$$

4.2 Comparison to an Auction with Discrete Signals

There is an analogy between this sequential search process and an auction procedure. The buyer is the auctioneer; the sellers are the bidders (thus the analogy here is to procurement auctions). The main differences from the usual auction scenario are: (i) auctioneer (buyer) cannot commit to a mechanism; (ii) auctioneer (buyer) incurs a cost per each bidder (seller) it samples; (iii) the sequentiality of the bids. These three differences matter in conjunction with one another. The sequentiality might not matter by itself if the buyer could commit. In a typical auction model the number of bidders is exogenously fixed whereas here it is determined endogenously and affected by the cost s.

In this section we compare the search scenario analyzed above to a standard auction scenario in the same setting. Suppose that the buyer holds a first (lowest) price (procurement) auction among n sellers. As above, each seller observes a signal $x \in \{\ell, h\}$ according the the same distributions. Then the sellers simultaneously submit bids and the buyer selects the lowest bid. Due to the discreteness of the signals, the symmetric equilibrium in this auction involves mixed strategies. Let us introduce the following notation:

M_x^n = the equilibrium distribution of prices offered by a seller with signal x when there are n sellers.

Q_w^n = the equilibrium price distribution offered by a random seller when there are n sellers and state is $w = f_{h|w} M_h^n + f_{\ell|w} M_\ell^n$.

$\underline{p}_x^n = \inf\{p : M_x^n(p) > 0\}$ and $\overline{p}_x^n = \sup\{p : M_x^n(p) < 1\}$.

$E(c \mid n \times h)$ = the expected cost when all n signals are h.

$\Pi_x^n(p)$ = the expected profit of a seller with signal x who offers p when there are n bidders and all the other follow their equilibrium strategies.

π_x^n = the expected equilibrium profit of a seller with signal x when there are n sellers. Thus, $\Pi_x^n(p) \leq \pi_x^n$, for all p and $\Pi_x^n(p) = \pi_x^n$ for all $p \in SUPP(M_x^n)$.

Proprosition 5: A symmetric equilibrium $M_x^n(p), x = \ell, h$, has the following properties.

(i) The support of $M_\ell^n(p)$ is an interval $[\underline{p}_\ell^n, \overline{p}_\ell^n)$ with no mass points; The support of $M_h^n(p)$ is an interval $[\underline{p}_h^n, \overline{p}_h^n]$ with a single mass point at $\overline{p}_h^n$.

(ii) $M_h^n(p) < M_\ell^n(p)$ for every $p < \overline{p}_h^n$.

(iii)

$$\begin{aligned}
\underline{p}_\ell^n &= \Pr(L \mid \ell)c_L + \Pr(H \mid \ell)c_H + \pi_\ell^n < \Pr(L \mid h)c_L + \Pr(H \mid h)c_H \\
&< \underline{p}_h^n \leq \overline{p}_\ell^n = \overline{p}_h^n = E(c \mid x^i = h, \forall i)
\end{aligned}$$

(iv) $\pi_\ell^n > 0 = \pi_h^n$. The total profit $n\pi_\ell^n \longrightarrow 0$ as $n \to \infty$.

4.3 Further Equilibria with Discrete Signals: Discussion of Restrictions

We restricted attention to a subclass of equilibria. Without this restriction, we can support different prices in the limit as outcomes of some equilibria. We do not have a strong view on which equilibria are the "right one." One might argue that the additional equilibria discussed below are somehow less robust. Independent of their robustness, the analysis of these equilibria provides further insights into the curiousities of search with adverse selection.

We will lift our restrictions one at a time, to separate their implications and to make it easier for the reader to judge whether the restriction is reasonable. The analysis is informal.

Non-Divinity. Suppose the equilibrium constellation σ is such that, at both sellers l and h, the prices are the same for both types of buyers and the trading probabilities are one,

$$M(l) = \left[p^l, 1, p^l, 1\right], \quad M(h) = \left[p^h, 1, p^h, 1\right].$$

Suppose $p^l < p^h$ and p^l and p^h are both strictly above expected costs, $p^x \geq E_0[c|x]$ but below c_H. The difference in prices is smaller than s. This mechanism is efficient and (therefore) undominated. The payoff of the L buyer is larger than the payoff of the H buyer, since the L buyer gets the low price p^l with higher probability, $f_{l|L} > f_{l|H}$. Suppose the posterior of the seller is $\beta(M, x) = 1$ for all off equilibrium mechanisms. Then, the mechanism satisfies all equilibrium definitions but it violates Divinity: The L buyer has a stronger incentive to offer a mechanism M' that specifies a lower price equal to expected costs, $M'(x) = [p', 1, p', 1]$. Fixing some $p_l > g_H c_H + g_L c_L$ and letting $p_h = p_l + s$, we can construct a sequence of equilibria with $s \to 0$, such that the limiting price is not equal to prior expected costs.

Non-monotone equilibria. The preferences of the buyer with respect to trading probabilities and prices depend on their continuation payoffs. If payoffs are ordered, these preferences satisfy the single crossing property such that the L buyer is the one preferring a low price. This property is usually assumed in applications but, in our model, preferences depend on an endogeneous object. Interestingly, we can construct equilibria in which preferences do not satisfy the single crossing property. Without the property, a belief refinement like Divinity, does not have much bite. Suppose now that σ is such that both buyers trade at the same price:

$$M(l) = [p, 1, p, 1], \quad M(h) = [p, 1, p, 1].$$

where p is larger than the expected costs of the h buyer and p is below c_H. The mechanism is again efficient and therefore undominated. Divinity has no bite since both buyers have the same continuation payoffs. Since they have the same continuation payoffs, they have the same preferences over prices and trading probabilities. Thefore, we can set $\beta(M) = 1$ for all off equilibrium mechanism. Thus, a constellation σ in which the mechanism specifies the same price for both buyers satisfies all equilibrium requirements but it fails monotonicity.

Dominated Mechanisms. There are equilibria in which dominated mechanisms are offered and in which the limiting prices are perfectly separating.

Remark: In an auction, with $N \to \infty$, the limit is never perfectly revealing. The reason we can construct perfectly revealing equilibria here is because the precision of information is endogeneous. In undominated equilibria this typically implies that signals become less informative. But it is illuminating for the richness of the model that one can find equilibria in which the opposite happens.

4.4 Continuous Signals

The above analysis uses the assumption that the signal can take on only two values (or, more importantly, a discrete number of values). However, some the main insights seem to be driven by the form of the winner's curse in the search environment, which does not seem as an artifact of the discrete space of signals. The objective of this section is to explore briefly the extension of the analysis to a richer signal space.

We are not going to repeat everything for richer signals. Instead we will focus here on the question of information revelation. Recall that, when the search cost s is small, no information is revealed in equilibrium in the sense that both types of the buyer end up transacting at the same price. It seems obvious that the discreteness of signals plays an important role in this result. We examine below how the extent of information revelation in equilibrium changes when there is a continuum of possible signals.

Suppose that the signal x can take on any value from $[0,1]$. As before $F_w(x)$ and $f_w(x)$ denote the distribution of and density of signals conditional on $w = L, H$ (Thus, $F_w(\cdot)$ has not atoms). Low values of x signal indicate a higher likelihood of L. That is, the likelihood ratio $\frac{f_L(x)}{f_H(x)}$ is decreasing in x. To simplify the setup, we assume that prior probabilities of the two states $\{L, H\}$ are $\frac{1}{2}$.

The question is to what extent is information revealed in equilibrium when s is small. The extent of revelation is captured here by the price paid by the L buyer when s is small. If the price that the L buyer pays is close to c_L and (therefore) the price that the H buyer pays is close to c_H, the revelation is maximal. Recall that the literature on auctions considered a related question. It inquired to what extent the equilibrium price in a common values auction reflects the correct information when the number of bidders is made arbitrarily large (Wilson(1977) and Milgrom(1979)). Milgrom's result translated to an auction version of our model is that the price approaches the true value iff $\lim_{x\to 0} \frac{f_L(x)}{f_H(x)} = \infty$. That is, when there are signals that are exceedingly more likely when the true state is L than when it is H. In our model the number of bidders is endogenous. The counterpart of increasing the number of bidders in our model is reduction of the sampling cost s. The following proposition claims that in our model revelation requires even stronger requirements on the quality of the signals.

To analyse the case of continuous signals, we make two normalizations that are without loss of generality. First, we set the low cost to zero, $c_L = 0$. Also, without loss of generality, signals are normalized such that the posterior probability of the high state is equal to x, i.e., for all x,

$$x = \frac{\frac{1}{2} f_H(x)}{\frac{1}{2} f_H(x) + \frac{1}{2} f_L(x)}.$$

Rewriting shows that this requires

$$\frac{f_H}{f_L} = \frac{x}{1-x}.$$

Therefore, the distribution $F_L(\cdot)$ determines $F_H(\cdot)$ and we can concentrate on characterizing $F_L(\cdot)$, the distribution of signals from the viewpoint of the $L-$Buyer.

Signals $x > 0$ are not revealing. As we have shown for discrete signals, if there are no revealing signals for the low state, the limit with $s \to 0$ involves complete pooling. The next lemma states that this is the case with a continuum of signals as well, i.e., this insights gained from the simpler discrete example is robust. (The proof is to be included.)

Theorem 3 *Suppose the suppport of F_L is $[a, b] \subset [0, 1]$. If $a > 0$, then the limit involves complete pooling at the ex ante expected price, i.e.,*

$$\bar{p}_L = \bar{p}_H = g_H c_H + g_L c_L.$$

In the example with discrete signals, the limit will be separating if there is a perfectly revealing signal for the low state. We ask now, whether the limit will be separating with a continuous signal distribution if its support includes zero, ie., if signals can be arbitrarily close to zero and therefore, signals can be arbitrarily informative. As noted before, in auctions it has been shown that the existence of such signals is sufficient for revelation of the state in the limit. As we will now see, this is not the case with search. The limit does not need to involve information revelation. Indeed, we will see that even with arbitrarily informative signals the limit can involve complete pooling. Thus, in a search model, the outcome can be very uninformative even in the presence of almost perfect information.

For the following analysis, it is helpful to analyse a special class of constellations, so called "canonical" constellations σ^E, first. A constellation is called a canonical constellation if both types of buyers offer the same mechanism $M^E(x)$, the mechanism M^E is accepted, reporting is truthful, and beliefs, given M^E, are equal to the prior β_0. Furthermore, a canonical constellation is described by a cutoff x^* and a probability d such that the mechanism $M^E(x)$ is equal to $[c_H, 0, c_H, d]$ if $x \geq x^*$ and $M^E(x) = [E_0[c|x], 1, E_0[c|x], 1]$ if $x \leq x^*$. So, for $x > x^*$ the L buyer does not trade and the H buyer trades at c_H with probability d. For $x \leq x^*$, both buyers trade with probability one at a price equal to the seller's expected cost. We require that the cutoff x^* is such that the L buyer is indifferent between trading at a price equal to expected costs and continuing search,

$$u - E_0\left[c|x^*, \sigma^E\right] = V_L\left(\sigma^E\right)$$

and we require that d is that the H trades if trading at c_H makes him strictly better off (and doesn't trade if trading at c_H makes him worse off)

$$d \begin{cases} = 1 & \text{if} \quad u - c_H > V_H\left(\sigma^E\right) \\ \in [0,1] & \text{if} \quad u - c_H = V_H\left(\sigma^E\right) \\ = 0 & \text{if} \quad u - c_H < V_H\left(\sigma^E\right). \end{cases}$$

A canonical constellation σ^E can be shown to exist. (The proof is to be included.)

Lemma 2 *An canonical constellation σ^E, described by (x^*, d) exists.*

A canonical constellation might or might not correspond to an equilibrium. The next lemma states that, if $V_H\left(\sigma^E\right) \leq u - c_H$, then there is actually an equlibrium σ^* that is outcome equivalent to σ^E (equivalent on the equilibrium path but not off the equilibrium path.) Conversely, if an equilibrium σ^* is such that $V_H(\sigma^*) \leq u - c_H$, then there is a canonical constellation σ^E that is equivalent to σ^*:

Lemma 3 *Suppose for some canonical constellation σ^E,$V_H\left(\sigma^E\right) \geq u - c_H$. Then there exists an equilibrium constellation σ^* such that, for all x, the mechanism offers are the same, $M^E(x) = M^*(x)$. Conversely, suppose for some equilibrium σ^*, $V_H(\sigma^*) \geq u - c_H$. Then there is some canonical constellation σ^E such that $M^*(x) = M^E(x)$.*

The tail of the distribution of signals. When $s_k \to 0$, buyers search for increasingly low signals close to zero. We therefore want to characterize the distribution of signals $x \leq x^*$ when $x^* \to 0$. We call the conditional distribution $\frac{F_L(x)}{F_L(x^*)}$ on $x \in [0, x^*]$ the *tail* of the distribution F_L. We concentrate on the set of distributions with "regular" tails, i.e., for $x^* \to 0$, the tail $\frac{F_L(x)}{F_L(x^*)}$ on $x \in (0, x^*]$

is well behaved in the following sense. First, for each x^*, we map the interval $(0, x^*]$ into $[0, \infty)$ via the continuous transformation

$$t(x, x^*) = \frac{x^* - x}{xx^*}$$

which defined as the solution to

$$x = \frac{x^*}{1 + x^* t(x, x^*)}.$$

So, $t(x^*, x^*) = 0$ and $\lim_{x \to 0} t(x, x^*) = \infty$. Therefore, for each x^*, $F_L(\cdot)$ induces a distribution $F_L^{x^*}$ on $[0, \infty)$ via

$$F_L^{x^*}(t) = 1 - \frac{F_L\left(\frac{x^*}{1+x^*t}\right)}{F_L(x^*)}.$$

For a sequence of cutoffs x^*, this mapping gives us a sequence of distributions of t on $[0, \infty)$. A tail is called **regular** if this distribution converges to a limit F_L^*. The set of distributions F_L which are regular is Φ,

$$\Phi = \left\{ \begin{array}{c} F_L(\cdot) : \exists\ F^*(t) \equiv 1 - \lim_{x^* \to 0} \frac{F_L\left(\frac{x^*}{1+x^*t}\right)}{F_L(x^*)} \in [0, 1], \ \forall t \\ \text{and } F^*(t) \leq 1 - \frac{F_L\left(\frac{x^*}{1+x^*t}\right)}{F_L(x^*)}, \forall t \in (0, \infty), \ \forall x^*. \end{array} \right\}$$

The second line of the definition is needed for the proof. It ensures that the integral of a linear function converges on an unbounded support, $\lim_{x^* \to 0} \int_0^\infty (at + c)\, dF_L^{x^*}(t) \to \int_0^\infty (at + c)\, dF^*(t)$ with a, c being some constants.

Every tail $F^*(t)$ for $F \in \Phi$ is exponential:

Lemma 4 *If $F \in \Phi$, then the limit tail $F^*(\cdot)$ is exponential, i.e., for some $\lambda \in [0, \infty]$,*

$$F^*(t) = 1 - e^{-\lambda t}.$$

The lemma is immediate if $F^*(t)$ is constant at 0 or constant at 1. (Note that we do not require F^* to be a proper cumulative distribution.) In these cases, $\lambda = 0$ and $\lambda = \infty$, respectively. If F^* is not constant, then it must have a stationarity property, since it must be independent of the cutoff x^*. This property requires that F^* is exponential. Therefore, we can define a mapping

$$\lambda : \Phi \to [0, \infty],$$

which assigns a hazard rate $\lambda(F_L)$ to each distribution $F_L \in \Phi$.

The next lemma characterizes the limiting prices of canonical constellation σ_k^E when $s_k \to 0$. Let $p(x_k^*)$ be the price offered to the cutoff seller, $p(x_k^*) = E_0\left[c | x^*, \sigma_k^E\right]$:

Lemma 5 *Consider a (sub-)sequence of economies along which $s_k \to 0$ and $p(x_k^*)$ converges to some $\bar{p}$. Then, for every sequence of canonical constellations $\{\sigma_k^E\}$*

$$\begin{array}{rcc} \bar{p} = c_L & \Leftrightarrow & \lambda(F_L) = 0 \\ \bar{p} = \lambda c_H + (1-\lambda) c_L & \Leftrightarrow & \lambda(F_L) \in \left(0, \frac{1}{2}\right) \\ \bar{p} = \frac{1}{2} c_H + \frac{1}{2} c_L & \Leftrightarrow & \lambda \geq \frac{1}{2}. \end{array}$$

Let $E_L[p|\sigma_k^*]$ be the expected price paid by the L buyer in expectation in a full equilibrium σ_k^*. We call the limit of a sequence of equilibria σ_k^* revealing if $E_L[p|\sigma_k^*] \to c_L$.

Theorem 4 ***(Main Result.)*** *Fix some distribution $F_L(\cdot) \in \Phi$ and some sequence $\{s_k\}$, $s_k \to 0$. Then, the limit of* every *sequence of corresponding equilibria $\{\sigma_k\}$ is revealing if and only if $\lambda(F_L) = 0$. If $\lambda(F_L) \geq \frac{1}{2}$, then a sequence of equilibria with complete pooling in the limit exists.*

Thus, revelation in the search model with small s requires that there are signals that separate L from H even in a more pronounced way than in the large auction model. When both models the signals that make L exceedingly more likely are needed to counteract the winner's curse. This difference between the strengths of the requirement in the two models owes to the somewhat different form of the winner's curse in these models. As explained before, in the search model the winner's curse is produced both by the larger expected number of sellers who participate in the bidding (like in the auction) and by the worsened distribution that a sampled seller is facing due to the longer search duration of the H type.[9]

[9] [Conclusion, Appendix and Literature to be added.]