Search and Adverse Selection*

Stephan Lauermann Asher Wolinsky

This version: April 2009

Abstract

This paper explores a dynamic model of adverse selection in which trading partners receive noisy information. A monopolistic buyer wants to procure service. Seller's cost depend on the buyer's type. The buyer contacts sellers sequentially and enters into a bilateral bargaining game. Each seller observes the buyer's offer. In addition, each seller observes a noisy signal. Contacting sellers (search) is costly. We characterize equilibrium when search cost become small. In the limit, the price will depend in a simple way on the curvature of the signal distribution. If signals are sufficiently strong, the limit outcome is equivalent to the full information outcome. (The equilibrium is separating and prices are equal to the true cost.) If signals are weak, the limit outcome is equivalent to an outcome with no information. (The equilibrium is pooling and prices are equal to ex ante expected cost.) The efficiency of the limit is closely tied to whether or not limit prices are separating or pooling. Intuitively, search cost reduce the winner's curse by reducing excessive search by bad types.

Away from the limit, a dynamic model of adverse selection with noisy information has several natural implications for the correlation between duration, quality, and prices. Most importantly, in many equilibria it will be the "lemons" that stay in the market for a long time, while good types trade fast. This is in accord with stylized facts about the housing or the labor market.

Very preliminary and Incomplete. Appendix not included

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

The paper looks at dynamic model of adverse selection. An agent that we call the buyer samples sequentially alternative trading partners, sellers, for a transaction that involves information asymmetry. The buyer knows his own characteristics (type), while sellers receive signals about it. The cost of the sellers depend on the buyer's characteristics. Signals are imperfect and the buyer has an incentive to search for a seller who received a favorable signal (indicating low cost). Sellers take this behavior into account when interpreting their own information. The main objective of this paper is to understand how the combination of search activity and information asymmetry affects prices and welfare.

Our main result concerns a situation in which search cost become small, i.e., we are looking at a limit. When search costs are small, equilibrium price can be characterized completely by the curvature of the tail of the signal distribution. We say that an equilibrium involves complete pooling if good and bad buyers trade at the same prices. An equilibrium is perfectly revealing if buyers get the same price they would get with perfect information. We show that, when search costs are small, equilibrium 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 or if the tail is not thick enough, the limit equilibrium involves complete pooling, i.e., prices are independent of individual characteristics. 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.

Whether or not information is perfectly revealed implies whether or not equilibrium is efficient. Except for signal distributions that are degenerate or have arbitrarily thin tails, equilibrium with small search cost is efficient if and only if the equilibrium is separating. We also discuss the relation between welfare and information revelation in an extension, where the efficient allocation depends on the personal characteristics of the buyer.

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] In contrast, in a model with search, the outcome can involve complete pooling even with arbitrarily informative signals. The main difference is that, with an auction, a buyer can commit to sample

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

[2] Note that we are looking at a model in which individual characteristics matter; Pesendorfer and Swinkels (2000) show that information aggregation is possible under weaker conditions for characteristics common to many buyers (e.g., the common value of stocks).

only a fixed number of sellers and buy from the seller with the lowest bid. In a search model, the buyer cannot commit to sample more sellers (or commit to truthfully report the number of sellers sampled before.)

Beyond auction theory, the relation between the strength of signals and information revelation by limit equilibria is analyzed in models of Herding and Voting. Duggan and Martinelli (2001) find that the existence of arbitrarily informative signals are a necessary and sufficient condition for information aggregation in a model of voting in juries when the size of the jury increases. Smith and Sorenson (2000) show that with social learning, herds on the correct action must occur if signals are arbitrarily informative (and weak conditions that that often suffice.)

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 and offers a *mechanism* which the seller can either accept or reject. If the mechanism is accepted, an allocation (trading probability and price) is implemented, depending on the reported type of the buyer. 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.[3] We show that the mechanism can be implemented as the outcome of a price proposal game between the buyer and the seller. Modelling the interaction as a mechanism proposal game has the advantage that we can concentrate ourself on pure strategy equilibria (thanks to the inscrutability principle). In a price proposal game, the price offer and acceptance strategies are generally mixed.

While the buyer can commit for the current period, he cannot commit not to trade in future periods. The buyer can also not provide evidence about the number of sellers already sampled, let alone provide evidence about their signals. (The buyer would have an incentive to commit to sample only a finite number of sellers and/or the buyer would like to truthfully communicate the number of sampled sellers, provided he has sampled only a few.) Equilibrium would be more efficient if the buyer could fully commit.

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.

More formally, we show that an appropriately transformed tail of the signal distribution can be approximated by an exponential distribtion function. Concentrating on the tail of the signal distribution for a low cost buyer, F_L, the parameter $\lambda \in [0, 1]$ of the approximating exponential distribution directly determines the equilibrium price. If $\lambda = 0$, (if the tail is thick), the low cost buyer will trade a price equal to cost c_L; If $\lambda \in \left(0, \frac{1}{2}\right)$, the limit price will be between c_L and ex ante expected cost; the limit price is strictly increasing in λ. If $\lambda \geq \frac{1}{2}$, the limit price is equal to prior expected cost.

We analyze the relation between information aggregation and welfare. In our base model, welfare is only affected by the accumulated search costs (buyers will purchase the good in every equilibrium and price are welfare irrelevant transfers). We show that accumulated search cost become zero in the limit, if the limit is separating (because then almost no bad buyer searches). However, accumulated search cost stay positive in all (partial) pooling equilibria, except when the tail of the signal distribution is arbitrarily thin.

[3] 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.

In an extension (not included in the current submission), we consider buyers with heterogeneous willingness to pay. In the efficient allocation, buyers with a low willingness to pay should receive service if and only if their type is good (only then is the cost of the service smaller than the valuation by the buyer). Whether or not the limit allocation is efficient depends on whether sellers can distinguish the types of the buyer. If the limit involves complete pooling (if signals are weak), the outcome is efficient.

Importantly, in this extension, smaller search cost can have a negative impact on welfare. With smaller search cost, bad buyers engage in more search and separation is harder to achieve. This is contrast with standart models of search in which smaller search cost increase welfare (directly) by increasing the match quality and (indirectly) by reducing the negative impact local market power.

In another extension (not included in the current submission), we consider a more structured search: a buyer first samples a small set of friends, before sampling strangers. Friends and strangers make different inference about the type of the buyer upon encounter. It takes stronger signal for a stranger to be willing to trade with a buyer than for a friend, i.e., strangers are more distrusting. We compare this result to the situation of an entrepreneur of a start-up company looking for an early investor. Convincing a friend (a member of an extended social network) to invest into a project seems much easier than convincing a stranger who is not socially connected to the entrepreneur.

We analyse limit equilibria for tractability. When search costs are not small, we cannot rule out multiplicity of equilibrium. For example, when good buyer sample more, sellers become more optimistic, making search more valuable. We discuss this in a separate section. We also illustrate the use of our refinements (for the principal agent game) in two lemmas following the main result. We show that we can get separating equilibria even without arbitrarily informative signals; however, such equilibria will involve (Pareto) dominated trading mechanisms. We also show that we can get pooling equilibria even if signals have a thick tail; however, such equilibria are supported by beliefs that fail devinity.

We discuss a numer of potential extensions. Most prominently, one can assume that the seller's signal is not observed by the buyer. As another extension, the buyer would learn his own type from either the signals or the rejection decisions by sellers (if the buyer does not observe sellers' signals.)

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[4]. 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 and engage in bargaining. 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. We also assume that $s > c_H - c_L$, otherwise search never pays.

At the beginning of every sampling round, the buyer draws one seller at a cost s. The seller receives a signal $x \in [a, b] \subset [0, 1]$ that is correlated with the state. The distribution of x given w is F_w. We assume that F_w is atomless and 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.[5] 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 $R = w$.[6] A typical mechanism will be

$$M = \left[c_L, q^C, c_H, 1\right],$$

which implies the following: If the seller accepts M and if the buyer reports a type H , then

[4]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.

[5]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.

[6]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.

trade happens at a price equal to c_H with probability one; if the buyer reports a type L, then trade happens at a price equal to c_L with probability $q^C \leq 1$. In general, $M(x, w)$ describes the mechanism offered by a buyer w if the seller has a signal x and a mechanism offer strategy is $M(\cdot, \cdot) : \{L, H\} \times X \to \mathbb{R}_+^4$.[7] Given a mechanism offer M, the seller believes that the probability of the high state is $\beta(M, x)$, so beliefs are $\beta(\cdot, \cdot) : X \times \mathbb{R}_+^4 \to [0, 1]$. $A(M, x)$ describes the acceptance decision by a seller of type x, $A(\cdot, \cdot) : \mathbb{R}_+^4 \times X \to [0, 1]$, 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, $R : \{L, H\} \times X \times \mathbb{R}_+^4 \to \{L, H\}$. We will define an equilibrium as a constellation in which strategies are mutually *optimal* and 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 accepted an offer M is equal to the expected profit from the contract conditional on the high cost and the low cost buyer, respectively, weighted by the relative probabilities

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

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. Let $\beta_0(x, \theta)$ denote the "interim" belief of a seller with signal

[7] 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.

x who is sampled by a buyer (β_0 does not condition on the price offer). The belief the seller depends on the relative prior likelihood of the types $\frac{g_L}{g_H}$, the relative likelihood of the the signals, $\frac{dF_L(x)}{dF_H(x)}$ and the relative likelihood of being sampled $\theta(\sigma)$. As shown in the appendix, the interim belief of the seller can be defined as

$$\beta_0(x,\theta) = \frac{1}{1 + \frac{g_L}{g_H}\frac{dF_L(x)}{dF_H(x)}\theta}.$$

(We implicitly assume an infinite number of sellers and the probability of being sampled follows an improper (uniform) prior. We derive β_0 as the limit of a Bayesian update with finitely many sellers when the number of sellers becomes large.) Note that the belief of a sampled seller depends on the equilibrium only through the ratio θ. In equilibrium, both buyers will offer the same mechanism and thus, the beliefs of a seller following an on-equilibrium mechanism offer are just $\beta(x,M) = \beta_0(x,\theta)$. 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.,

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

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 discuss section.

Beliefs following an off equilibrium mechanism offer M' satisfy "*Divinity*" if they put (weakly) higher probability on a type of buyer who is strictly better off if M' is accepted (rather than trading at the equilibrium mechansim). With $U_w^*(x)$ denoting the equilibrium payoff, let $U_w(M', x)$ be the payoff to buyer w if the mechanism M' is accepted,

$$U_w\left(M', x\right) = q_R^{M'}\left(u - p_R^{M'}\right) + \left(1 - q_R^{M'}\right)V_w - s,$$

given optimal reporting. Beliefs $\beta(x, M')$ satisfy divinity given σ if

$$\begin{aligned} \beta\left(x, M'\right) &\geq \beta_0(x, \theta) \text{ if } U_H\left(M', x\right) > U_w^*(x) \\ \beta\left(x, M'\right) &\leq \beta_0(x, \theta) \text{ if } U_L\left(M', x\right) > U_L^*(x) \end{aligned}$$

Divinity (rather than refinements like D1/D2) is used because it makes the construction of equilibrium easier; for example, assigning the belief $\beta_0(x, \theta)$ off the equilibrium path would ensure that an equilibrium satisfies Divinity. Note that $\beta(x, M') = \beta_0(x, \theta)$ whenever both buyers strictly prefer M' to the equilibrium mechanism.

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]

Divinity in the current definition implies that equilibrium mechnisms must be *undominated*. We state this as an extra requirement for transparency. A mechanism M is undominated if there is no other mechansim M'' such that both types of the buyer seller strictly prefer M'' to M and seller's expected profits under M' are strictly higher than under M, given the interim belief β_0. (Of course, if both types of buyers strictly prefer a mechanism M'' to M, then divinity requires

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

that the belief of the seller is equal to β_0. If seller's profits are positive, sequential rationality requires him to accept a mechanism. Therefore, divinity implies that equilibrium mechanisms must be undominated.)

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. We also characterize the set of mechanisms $M(x)$ that satisfy the equilibrium definition for given continuation payoffs $V_w(\sigma)$ and interim beliefs $\beta_0(x, \theta(\sigma))$.

Given a ratio θ, the expected cost of a seller with signal x is

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

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

We show that in any equilibrium, the mechanism that is offered in any given buyer-seller pair must maximizes the payoff of the L buyer, subject to feasibility constraints (the mechanism should be weakly profitable for the seller, reporting should be truthful, and the H buyer should not prefer to reveal his type and trade at a price equal to high cost c_H). Furthermore, every equilibrium is equivalent (in terms of expected prices, number of expected searches and payoffs) to an equilibrium that is characterized by three numbers: $\mathbf{x} = \left[x^*, x^{**}, q_L^C\right]$. With $Ec = E_0[c|x, \theta]$ denoting the interim expected cost of a seller, the mechanism that is offered to a seller with signal x is given by

$$M(x) = \begin{cases} [1, \quad Ec, \quad 1, \quad Ec] & \text{If} \quad x \leq x^* \\ \left[q_L^C, \quad c_L, \quad 1, \quad c_H\right] & \text{if} \quad x \in (x^*, x^{**}) \\ [\,0, \quad c_L, \quad 0, \quad c_H] & \text{if } x \geq x^{**}. \end{cases} \tag{1}$$

Thus, if the signal is low, $x \in [0, x^*]$, both buyers trade at the same price equal to the interim expected cost. If the signal is high, $x \geq x^{**}$, the trading probability is zero. The trading probability is positive if the signal is intermediate, $x \in (x^*, x^{**})$, but, of course, the trading probability at the lower price c_L cannot be one; otherwise, the mechanism would not be incentive compatible. Instead, q_L^C will make the H buyer just indifferent between trading at c_H with probability one and trading at c_L with probability q_L^C,

$$q_L^C\left(V_H, V_L, \theta\right) = \begin{cases} \frac{u - c_H - V_H}{u - c_L - V_H} & \text{if} \quad u - c_H - V_H > 0 \\ 0 & \text{if } u - c_H - V_H \leq 0 \end{cases}$$

The trading probability at c_L in the intermediate region (x^*, x^{**}) is positive only if $u - c_H - V_H > 0$; otherwise, it is not possible to make the H buyer indifferent. The cutoff x^* is always strictly positive while x^{**} can be one. The cutoff x^* corresponds to a signal such that the L buyer is indifferent between trading at a price equal to the expected cost of the seller and trading at the price c_L with probability q^C:

$$x^*\left(V_H, V_L, \theta\right) : u - E_0\left[c|x^*, \theta\right] = q_L^C\left(u - c_L\right) + \left(1 - q_L^C\right) V_L,$$

The cutoff x^{**} can be anything in $[x^*, 1]$.

The next lemma states that every equilibrium is equivalent to one in which the mechanism is as described before:

Lemma 1 *Given any equilibrium σ^* with payoffs $V_H(\sigma^*)$ and $V_L(\sigma^*)$, and ratio $\theta(\sigma^*)$. Then there is an equilibrium σ^{**} in which the offered mechanism is described by some $\mathbf{x} = \left[x^*, x^{**}, q_L^C\right]$, with $x^* = x^*\left(V_H, V_L, \theta\right)$, $x^{**} \geq x^*$, and $q_L^C = q_L^C\left(V_H, V_L, \theta\right)$ such that with $Ec = E_0\left[c|x, \theta\right]$*

$$M^{**}(x) = \begin{cases} [1, \quad Ec, \quad 1, \quad Ec] & \textit{If} \quad x < x^* \\ \left[q_L^C, \quad c_L, \quad 1, \quad c_H\right] & \textit{if} \quad x \in (x^*, x^{**}) \\ [\,0, \quad c_L, \quad 0, \quad c_H] & \textit{if } x > x^{**}. \end{cases}$$

And σ^{**} leads to the same payoffs and ratio as σ^*.

We also show that equilibrium exists:

Theorem 1 *Equilibrium exists.*

4 Main Result

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, 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 $\lim_{x\to a} \frac{f_L(x)}{f_H(x)} = \infty$. 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. In addition, sellers' posteriors can be expressed very as a function of the signal and the relative number of searches as shown below. For further simplification, we often set $c_H = 1$, so that expected cost of a seller are equal to the belief β.

Signals $x > 0$ are not revealing. If there are no revealing signals for the low state, the limit with $s \to 0$ involves complete pooling. Let $E_L[p|\sigma_k]$ be the expected price paid by the L buyer in expectation in an equilibrium σ_k, given s_k. We say that the limit of a sequence of equilibria σ_k^* involves complete pooling if $E_L[p|\sigma_k^*] \to g_H c_H + g_L c_L$.

Theorem 2 *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.,$E_L[p|\sigma_k^*] \to g_H c_H + g_L c_L$.*

Proof: Take a sequence of constellations σ_k for $s_k \to 0$. Let $V_{kw} = V_w(\sigma_k)$ and $\Delta_k = u - c_H - V_{kH}$. In general, a subscript k denotes parameters of the constellation σ_k (like x_k, θ_k, etc.). We distinguish three cases according to whether or not Δ_k is positive, zero, or negative when k is large (if the sign of Δ_k does not converge, the analyis is for an arbitrary convergent subsequence which is sufficient for the conclusion). We will only consider the first case, $\Delta_k < 0$, here. The other cases are appendicized. For k large enough, in equilibrium $\Delta_k < 0$.

Case 1: $\Delta_k < 0$ for all k large enough. Then $x_k^{**} = x_k^*$ and both buyers search for a seller with a signal $x \leq x_k^*$. The ratio of the number of searches is

$$\theta_k = \frac{F_H\left(x_k^*\right)}{F_L\left(x_k^*\right)}.$$

The cutoff x_k^* is determined by indifference of the L buyer between trading at the expected cost of a seller with this signal and continuing search

$$x_k^* : E_0\left[c|x_k^*, \theta_k\right] - \int_a^{x_k^*} E_0\left[c|x, \theta_k\right] \frac{dF_L(x)}{F_L\left(x_k^*\right)} = \frac{s_k}{F_L\left(x_k^*\right)}.$$

The cutoff x_k^* must converge to the lower bound of the support, a. Otherwise, search cost on the right hand side converge to zero, while the expected saving from search on the left hand side would be positive: Since x_k^* is bounded away from a, the ratio θ_k is bounded away from the extremes, 0 and ∞. Hence, sellers with different signals will offer different price.

Let $x_k^* \to a$. Then the ratio becomes equal to the inverse likelihood ratio,

$$\begin{aligned}
\lim \theta_k &= \lim \frac{F_H\left(x_k^*\right)}{F_L\left(x_k^*\right)} \\
&= \lim \frac{f_H\left(x_k^*\right)}{f_L\left(x_k^*\right)} = \frac{a}{1-a}.
\end{aligned}$$

The expected price at which the L buyer is trading is

$$\begin{aligned}
\lim E_L\left[p|\sigma_k^*\right] &= \lim \int_a^{x_k^*} \frac{1}{1+\frac{1-x}{x}\theta_k} c_H \frac{dF_L(x)}{F_L\left(x_k^*\right)} \\
&= \frac{1}{1+\frac{1-a}{a}\frac{a}{1-a}} c_H = \frac{1}{2} c_H = E_0 c.
\end{aligned}$$

Hence, if $\Delta_k < 0$ for all k, the limit involves complete pooling. QED

The intuition is this: When search cost are small, both buyers search for sellers with the most favorable signals close to the lower bound a. The resulting ratio of the number of searches is

$$\lim \theta_k = \frac{a}{1-a}.$$

Of course, this is just the inverse likelihood ratio of the signals

$$\lim \frac{f_L\left(x_k^*\right)}{f_H\left(x_k^*\right)} = \frac{1-a}{a}.$$

Intuitively, if a high cost buyer is less likely to generate a signal close to the lower bound a, high cost buyers are searching even more. Hence, the informational content of the signals at the lower bound is just balanced by the informational content of being sampled.

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.

The intuition is this: If the limit is separating, the L buyer trades at c_L while the H buyer trades at c_H. It can be shown that the accumulated search cost of the L buyer must become zero. Hence, the L buyer must be able to find prices close to c_L at almost no cost. However, the search cost for the H buyer must be strictly positive. As we have seen before, this is not possible if the support of the signal distribution is bounded away from zero. Our main result shows that something similar happens when the support of the signal distribution is too thin near zero.

We will first look at equilibria in which the surplus of the H buyer is non-positive, $\Delta_k \leq 0$. In such equilibria, the cutoff x_k^* must converge to zero. This is intuitive: The L buyer can otherwise search for signals x close to zero, ensuring trade at a price close to c_L at almost no cost.

The incentive of the L buyer will depend strongly on the shape of the conditional distribution $\frac{dF_L(x)}{F_L(x_k^*)}$ on the left tail $(0, x_k^*)$. If this conditional distribution has a "thick" tail and puts a high mass on signals strictly below x_k^*, search will be more valuable (because the average seller will offer a strictly better deal) than in the case of a "thin" tail, when the conditional distribtion puts high mass on signals very close to x_k^* itself (because the average seller $x \leq x_k^*$ will offer almost the same deal as x_k^*). We therefore introduce a way of characterising the limiting tail distribution.

Given a constellation σ_k, the cutoff x_k^* is determined by indifference of the L buyer

$$E_0\left[c|x_k^*, \theta_k\right] - E_0\left[c|x \leq x_k^*, \theta_k, L\right] = \frac{s_k}{F_L\left(x_k^*\right)}.$$

The assumption $\Delta_k \leq 0$ implies that the H buyer has a weak incentive to not trade at c_H but rather incure search cost and find some $x \leq x_k^*$,

$$I\left(x_k^*, \theta_k\right) = c_H - E_0\left[c|x \leq x_k^*, \theta_k, H\right] - \frac{s_k}{F_H\left(x_k^*\right)} \geq 0.$$

We can use the indifference condition of the L buyer to substitute s_k. Using that $dF_H(x) = \frac{x}{1-x} dF_L(x)$, we get (see Appendix):

$$I(x_k^*, \theta_k) F_H(x_k^*) = \int_0^{x_k^*} \left(\left(1 - \frac{1}{1 + \frac{(1-x)}{x}\theta_k} \right) \frac{x}{1-x} - \frac{1}{1 + \frac{(1-x_k^*)}{x_k^*}\theta_k} + \frac{1}{1 + \frac{(1-x)}{x}\theta_k} \right) dF_L(x)$$

We rewrite this expression further. Let C_k be the likelihood ratio $\frac{f_L(x_k^*)}{f_H(x_k^*)}\theta_k$ at the cutoff seller,

$$E_0 [c|x_k^*, \theta_k] = \frac{1}{1 + \frac{dF_L(x^*)}{dF_H(x^*)}\theta} = \frac{1}{1 + C_k}.$$

The price at the cutoff seller converges to prior expected cost if $C_k \to 1$. The price converges to $c_L = 0$ if $C_k \to \infty$. If $C_k \to \bar{C} \in (1, \infty)$, the limit price at the cutoff seller is in between.

We also do a change of variables. For each x_k^*, we map the interval $(0, x_k^*]$ into $[0, \infty)$ via the continuous transformation

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

which defined as the solution to $x = \frac{x_k^*}{1 + x_k^* t(x, x_k^*)}$. 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

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

Substituting into I_k we get

$$I(x_k^*, \theta_k) = \frac{x^* F_L(x_k^*)}{F_H(x_k^*)} \int_0^\infty \left(\frac{C_k}{1 - x_k^* + (1 + x^*(t-1)) C_k} - \frac{tC_k}{(1 + C_k)\left(1 - x_k^* + (1 + x^*(t-1)) C_k\right)} \right) f^{x_k^*}(t).$$

When $x_k^* \to 0$, we can evaluate the sign of the limit of $I(x_k^*, \theta_k)$ if we can pass the limit into the integral. As we will discuss now, if $f^{x_k^*}(t)$ converges, it must converge to an exponential with parameter λ. This will allow us to characterise the incentives of the buyer by the parameter λ.

More precisely, given a sequence $x_k^* \to 0$, we have a sequence of distributions $F_L^{x_k^*}(t)$ of t on $[0, \infty)$. Each of these distributions corresponds to a distributions $\frac{F_L(x(t))}{F_L(x_k^*)}$ of signals $x(t)$ on the tail $[0, x_k^*]$. A tail $\frac{F_L(x(t))}{F_L(x_k^*)}$ is called **regular** if the corresponding distribution $F_L^{x_k^*}(t)$ converges to a limit F_L^*. (The limit F_L^* does not need to be a cumulative distribution function itself.) The set of distributions F_L which are regular is Φ,

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

Note that by Helley's selection theorem, all distribution functions have a pointwise convergent subsequence, so a limit as in line one exists for all distributions. The second line of the definition is a technical condition which is needed for the proof. It ensures that the integral of a linear function converges. A generic example of functions $F_L \in \Phi$ is $F_L(x) = e^{-k\frac{1}{x}+k}$, since $\lambda(F_L) = k$ and

$$F_L^{x_k^*}(t) = 1 - \frac{e^{-k\left(\frac{x^*}{1+x^*t}\right)^{-1}+k}}{e^{-k\frac{1}{x^*}+k}} = 1 - e^{-kt} \quad \forall x_k^*.$$

Every tail $F^*(t)$ for $F \in \Phi$ is exponential: (The appendix contains the proof.)

Lemma 2 *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 the limit $F^*(t)$ is constant at 0 or constant at 1. In these cases, $\lambda = 0$ and $\lambda = \infty$, respectively. Many distributions will have such a degenerate limit. 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 $F_L \in \Phi$.

Using the substitions from before, the sign of I in the limit is the sign of

$$\begin{aligned} & \int_0^\infty \lim \left(\frac{C_k}{1 - x_k^* + (1 + x^*(t-1))\,C_k} - \frac{tC_k}{(1+C_k)\left(1 - x_k^* + (1 + x^*(t-1))\,C_k\right)} \right) \lim f^{x_k^*}(t) \\ = \ & \frac{\bar{C}}{\left(1+\bar{C}\right)^2} \int_0^\infty \left(1 + \bar{C} - t\right) \lambda e^{-\lambda t} dt \\ = \ & \frac{\bar{C}}{\left(1+\bar{C}\right)^2} \left(1 + \bar{C} - \frac{1}{\lambda}\right). \end{aligned}$$

(In the appendix we discuss why we can pass the limit into the integral given our assumptions on F_L.)

The H buyer has an incentive to search if and only if $I_k \geq 0$. We can use this to characterize equilibria with $\Delta_k \leq 0$. (To characterize equilibria $\Delta_k > 0$, we need to take care of q_k^C which makes the limit expressions more complicated.) Under which conditions $\Delta_k < 0$ for all k large enough? If $\Delta_k < 0$, the equilibrium must involve complete pooling in the limit: both buyers search and the expected price must be equal to the prior expected price (the search cost of the L buyer converge to zero; hence, the expected price conditional on $x \leq x_k^*$ must be equal to the price at the cutoff type for indifference.) Hence $C_k \to \bar{C} = 1$. Furthermore, $\Delta_k < 0$ requires $I(x_k^*, \theta_k) > 0$ for all k. Inspecting the limit expression shows that his is the case only if

$$1 + 1 - \frac{1}{\lambda} \geq 0 \Leftrightarrow \lambda \geq \frac{1}{2}.$$

Hence, we will get an equilibria with $\Delta_k < 0$ for all k only if $\lambda \geq \frac{1}{2}$.

Now, under which conditions $\Delta_k = 0$ for all k large enough? $\Delta_k = 0$ requires that

$$1 + \bar{C} - \frac{1}{\lambda} = 0.$$

Hence, $\Delta_k = 0$ for k large only if the limit price is $\bar{p} = \frac{1}{1+\bar{C}} = \lambda$. And hence, $\Delta_k = 0$ only if $\lambda \leq \frac{1}{2}$ (otherwise, $\bar{p} > \frac{1}{2}$, which contradicts seller's zero profits.)

In the appendix we show that $\Delta_k > 0$ for all k only if $\lambda = 0$. Of course, $\Delta_k > 0$ implies that the H buyer does not search while the L buyer searches for a seller with a signal close to zero. Hence, the limit must be revealing, $C_k \to \infty$.

We can characterize equilibrium prices by λ. Let $E_L[p|\sigma_k]$ be the expected price paid by the L buyer in expectation in an equilibrium σ_k. We call the limit of a sequence of equilibria σ_k^* revealing if $E_L[p|\sigma_k^*] \to c_L$. Recall, $c_L = 0$ and $c_H = 1$, and $prob\{w = H\} = prob\{w = L\} = \frac{1}{2}$.

Theorem 3 ***(Main Result.)*** *Fix some distribution $F_L \in \Phi$ and some sequence $\{s_k\}$, $s_k \to 0$. Let σ_k^* be a sequence of equilibria given s_k. Then the limit price paid by the L buyer is*

$$\lim E_L[p|\sigma_k^*] = \begin{cases} 0 & \text{if} \quad \lambda(F_L) = 0 \\ \lambda & \text{if} \quad \lambda(F_L) \in \left(0, \frac{1}{2}\right) \\ \frac{1}{2} & \text{if} \quad \lambda(F_L) \geq \frac{1}{2}. \end{cases}$$

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]

Furthermore:

Theorem 4 *The limit is efficient, $g_H V_H + g_L V_L = g_H (u - c_H) + g_L (u - c_L)$, if and only if either $\lambda (F_L) = 0$ or $\lambda (F_L) = \infty$.*

Thus, the limit is only efficient if either the limit is separating or signals are extremely weak. (An earlier example of an efficient limit with weak signals was the case with $a > 0$.)

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