

Screening and Adverse Selection in Frictional Markets*

Benjamin Lester
Federal Reserve Bank of Philadelphia

Ali Shourideh
Carnegie Mellon University

Venky Venkateswaran
NYU – Stern School of Business

Ariel Zetlin-Jones
Carnegie Mellon University

October 5, 2017

Abstract

We incorporate a search-theoretic model of imperfect competition into a standard model of asymmetric information with unrestricted contracts. We characterize the unique equilibrium, and use our characterization to explore the interaction between adverse selection, screening, and imperfect competition. We show that the relationship between an agent's type, the quantity he trades, and the price he pays is jointly determined by the severity of adverse selection and the concentration of market power. Therefore, quantifying the effects of adverse selection requires controlling for market structure. We also show that increasing competition and reducing informational asymmetries can decrease welfare.

Keywords: Adverse Selection, Imperfect Competition, Screening, Transparency, Search Theory
JEL Codes: D41, D42, D43, D82, D83, D86, L13

*We thank Gadi Barlevy, Hanming Fang, Mike Golosov, Piero Gottardi, Veronica Guerrieri, Ali Hortacsu, Alessandro Lizzeri, Guido Menzio, Derek Stacey, Robert Townsend, Randy Wright, and Pierre Yared, along with seminar participants at the Spring 2015 Search and Matching Workshop, 2015 CIGS Conference on Macroeconomic Theory and Policy, 2015 annual meeting of the Society for Economic Dynamics, 2015 meeting of the Society for the Advancement of Economic Theory, 2015 Norges Bank Conference on Financial Stability, 4th Rome Junior Conference on Macroeconomics (EIEF), 2015 Summer Workshop on Money, Banking, Payments and Finance at the Federal Reserve Bank of St. Louis, 2015 West Coast Search and Matching, 2015 Vienna Macro Workshop, NYU Search Theory Workshop, European University Institute, Toulouse School of Economics, University of Pennsylvania, Columbia University, the Federal Reserve Bank of Cleveland, and the Banque de France for useful discussions and comments. The views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. This paper is available free of charge at philadelphiafed.org/research-and-data/publications/working-papers.

1 Introduction

Many important markets suffer from adverse selection, including the markets for insurance, credit, and certain financial securities. There is mounting evidence that many of these markets also feature some degree of imperfect competition.¹ And yet, perhaps surprisingly, the effect of imperfect competition on prices, allocations, and welfare in markets with adverse selection remains an open question.

Answering this question is important for several reasons. For one, many empirical studies attempt to quantify the effects of adverse selection in the markets mentioned above.² A natural question is to what extent these estimates—and the conclusions that follow—are sensitive to the assumptions imposed on the market structure. There has also been a recent push by policymakers to make these markets more competitive and less opaque.³ Again, a crucial, but underexplored question is whether these attempts to promote competition and reduce information asymmetries are necessarily welfare-improving.

Unfortunately, the ability to answer these questions has been constrained by a shortage of appropriate theoretical frameworks.⁴ A key challenge is to incorporate nonlinear pricing schedules—which are routinely used to screen different types of agents—into a model with asymmetric information and imperfect competition. This paper delivers such a model: we develop a novel, tractable framework of adverse selection, screening, and imperfect competition.

The key innovation is to introduce a search-theoretic model of imperfect competition (à la ?) into an otherwise standard model with asymmetric information and nonlinear contracts. Within this environment, we provide a full analytical characterization of the unique equilibrium, and then use this characterization to study both the positive and normative issues highlighted above.

First, we show how the structure of equilibrium contracts—and hence the relationship between an agent’s type, the quantity that he trades, and the corresponding price—is *jointly* determined by the severity of the adverse selection problem and the degree of imperfect competition. In particular, we show that equilibrium offers separate different types of agents when competition is relatively intense or adverse selection is relatively severe, while they typically pool different types of agents in markets where principals have sufficient market power and adverse selection is sufficiently mild. Second, we explore

¹For evidence of market power in insurance markets, see ?, ?, and ?. For evidence of market power in various credit markets, see, e.g., ?, ?, and ?. In over-the-counter financial markets, a variety of data suggest that dealers extract significant rents; indeed, this finding is hard-wired into workhorse models of this market, such as ?.

²See the seminal paper by ?, and ? for a comprehensive survey.

³Increasing competition and transparency in health insurance markets is a cornerstone of the Affordable Care Act, while the Dodd-Frank legislation addresses similar issues in over-the-counter financial markets. In credit markets, on the other hand, legislation has recently focused on *restricting* how much information lenders can demand or use from borrowers.

⁴As ? put it, “there is a crying need for [a model] devoted to the interaction between imperfect competition and adverse selection.”

how total trading volume—which, in our environment, corresponds to the utilitarian welfare measure—responds to changes in the degree of competition and the severity of adverse selection. We show that increasing competition or reducing informational asymmetries is only welfare-improving in markets in which *both* market power is sufficiently concentrated and adverse selection is sufficiently severe.

Before expanding on these results, it is helpful to lay out the basic ingredients of the model. The agents, whom we call “sellers,” are endowed with a perfectly divisible good of either low or high quality, which is private information. The principals, whom we call “buyers,” offer menus containing price-quantity combinations to potentially screen high- and low-quality sellers.⁵ Sellers can accept at most one contract, i.e., contracts are exclusive. To this otherwise canonical model of trade under asymmetric information, we introduce imperfect competition by endowing the buyers with some degree of market power. The crucial assumption is that each seller receives a stochastic number of offers, with a positive probability of receiving only one. Hence, when a buyer formulates an offer, he understands that it will be compared with an alternative offer with some probability, which we denote by π , and it will be the seller’s only option with probability $1 - \pi$. This formulation allows us to capture the perfectly competitive case by setting $\pi = 1$, the monopsony case by setting $\pi = 0$, and everything in between.

For the general case of imperfect competition, with $\pi \in (0, 1)$, the equilibrium involves buyers mixing over menus according to a nondegenerate distribution function.⁶ Since each menu is comprised of two price-quantity pairs (one for each type), the main equilibrium object is a probability distribution over four-dimensional offers. An important contribution of our paper is developing a methodology that allows for a complete, yet tractable, characterization of this complicated equilibrium object.

We begin by showing that any menu can be summarized by the indirect utilities it offers to sellers of each type. Next, we establish an important property: in any equilibrium, all menus that are offered by buyers are ranked in exactly the same way by both low- and high-quality sellers. This property, which we call “strictly rank-preserving,” implies that all equilibrium menus can be ranked along a single dimension. The equilibrium, then, can be described by a distribution function over a unidimensional variable—say, the indirect utility offered to low-quality sellers—along with a strictly monotonic function mapping this variable to the indirect utility offered to the high-quality seller. We show how to solve for

⁵The use of the labels “buyers” and “sellers” is merely for concreteness and corresponds most clearly with an asset market interpretation. These monikers can simply be switched in the context of an insurance market, so that the “buyers” of insurance are the agents with private information and the “sellers” of insurance are the principals.

⁶Mixing is to be expected for at least two reasons. First, this is a robust feature of nearly all models in which buyers are both monopsonists and Bertrand competitors with some probability, even without adverse selection or nonlinear contracts. Second, even in perfectly competitive markets, it is well known that pure strategy equilibria may not exist in an environment with both adverse selection and nonlinear contracts.

these two functions, obtaining a full analytical characterization of all equilibrium objects of interest, and then establish that the equilibrium is unique. Interestingly, our approach not only avoids the well-known problems with existence of equilibria in models of adverse selection and screening, but also requires no assumptions on off-path beliefs to get uniqueness. We then use this characterization to explore the implications of imperfect competition in markets suffering from adverse selection.

First, we show that the structure of menus offered in equilibrium depends on both the degree of competition, captured by π , and the severity of the adverse selection problem, which is succinctly summarized by a single statistic that is largest (i.e., adverse selection is most severe) when: (i) the fraction of low-quality sellers is large; (ii) the potential surplus from trading with high-quality sellers is small; and (iii) the information cost of separating the two types of sellers, as captured by the difference in their reservation values, is large. Given these summary statistics, we show that separating menus are more prevalent when competition is relatively strong or when adverse selection is relatively severe, while pooling menus are more prevalent when competition is relatively weak and adverse selection is relatively mild. Interestingly, holding constant the severity of adverse selection, the equilibrium may involve all pooling menus, all separating menus, or a mixture of the two, depending on the degree of competition. This finding suggests that attempts to infer the severity of adverse selection from the distribution of contracts that are traded should take into account the extent to which the market is competitive.

Next, we examine our model's implications for welfare, defined as the objective of a utilitarian social planner. In our context, this objective maps one-for-one to the expected quantity of high-quality goods traded. A key finding is that competition can worsen the distortions related to asymmetric information and, therefore, can be detrimental to welfare. When adverse selection is mild, these negative effects are particularly stark: welfare is actually (weakly) maximized under monopsony, or $\pi = 0$.

When adverse selection is severe, however, welfare is inverse U-shaped in π , i.e., an *interior* level of competition maximizes welfare. To understand why, note that an increase in competition induces buyers to allocate more of the surplus to sellers (of both types) in an attempt to retain market share. All else equal, increasing the utility offered to low-quality sellers is good for welfare: by relaxing the low-quality seller's incentive compatibility constraint, the buyer is able to exchange a larger quantity with high-quality sellers. However, *ceteris paribus*, increasing the utility offered to high-quality sellers is bad for welfare: it tightens the incentive constraint and forces buyers to trade less with high-quality sellers. Hence, the net effect of an increase in competition depends on whether the share of the surplus offered to high-quality sellers rises faster or slower than that offered to low-quality sellers.

When competition is low, buyers earn a disproportionate fraction of their profits from low-quality sellers. Therefore, when buyers have lots of market power, an increase in competition leads to a faster increase in the utility offered to low-quality sellers, since buyers care relatively more about retaining these sellers. As a result, the quantity traded with high-quality sellers and welfare rise with competition. When competition is sufficiently high, profits come disproportionately from high-quality sellers. In this case, increasing competition induces a faster increase in the utility offered to high-quality sellers and, therefore, a decrease in expected trade and welfare. These results suggest that promoting competition—or policies that have similar effects, such as price supports or minimum quantity restrictions—can have adverse effects on welfare in markets that are sufficiently competitive and face severe adverse selection.

Next, we study the welfare effects of providing buyers with more information—specifically, a noisy signal—about the seller’s type. As in the case of increasing competition, the welfare effects of this perturbation depend on the severity of the two main frictions in the model: imperfect competition and adverse selection. When adverse selection is relatively mild or competition relatively strong, reducing informational asymmetries can actually be detrimental to welfare. The opposite is true when adverse selection and trading frictions are relatively severe. In sum, these normative results highlight how the interaction between these two frictions can have surprising implications for changes in policy (or technological innovations), underscoring the need for a theoretical framework such as ours.

Our baseline model, which we describe in Section 2 and analyze in Sections 3–5, was designed to be as simple as possible in order to focus on the novel interactions between adverse selection and imperfect competition. In Sections 6 and 7, we analyze several relevant extensions and variants of our model. In Section 6, we endogenize the level of competition by letting buyers choose the intensity with which they “advertise” their offers. This allows us to study how the severity of adverse selection can influence the market structure, and the ensuing welfare implications. In Section 7, we consider a more general market setting with an arbitrary meeting technology, where sellers can meet any number of buyers (including zero). We show how to derive the equilibrium in this setting, using the techniques from our benchmark model, and confirm that our main welfare results hold for certain popular meeting technologies.⁷

Literature Review. Our paper contributes to the extensive body of literature on adverse selection and, specifically, the role of contracts as screening devices. Most of this literature has either assumed a

⁷In the Appendix, we explore a number of additional extensions: we relax the assumption of linear utility to analyze the canonical model of insurance under private information; we allow the degree of competition to differ across sellers of different quality; we show how to incorporate additional dimensions of heterogeneity, including horizontal and vertical differentiation; and we consider the case of $N > 2$ types or qualities.

monopolistic market structure (a la ?) or perfect competition (a la ?).⁸ The main novelty of our analysis is to synthesize a standard model of adverse selection and screening with the search-theoretic model of imperfect competition developed by ?. While this model of imperfect competition has been used extensively in both theoretical and empirical work, to the best of our knowledge none of these papers address adverse selection and screening.⁹

A recent paper by ? exploits the ? model in an environment with screening contracts and asymmetric information, but the asymmetric information is over the agents' *private values*. This key difference implies that the role of screening—and how it interacts with imperfect competition—is ultimately very different in our paper and theirs.¹⁰

More closely related to our work is the literature that studies adverse selection and nonlinear contracts in an environment with *competitive search*, such as ?.¹¹ As in our paper, ? present an explicit model of bilateral trade without placing any restrictions on contracts, beyond those arising from the primitive frictions. There are, however, several important differences. First, we study how perturbations to the search technology affect market power, and the interaction between the resulting distortions and the underlying adverse selection problem, while ? and others focus on the role of search frictions in providing incentives (through the probability of trade) and not on market power per se. Second, depending on parameters, our equilibrium menus can be pooling, separating, or a combination of both; the equilibrium in ? always features separating equilibria. In this sense, our approach has the potential to speak to a richer set of observed outcomes. Finally, we obtain a unique equilibrium without additional assumptions or refinements, whereas uniqueness in ? relies on a restriction on off-equilibrium beliefs.

An alternative approach to modeling imperfect competition is through product differentiation, as in ?, ?, ?, ?, and ?.¹² These papers vary competition by perturbing the importance of an orthogonal attribute of each contract, which is interpreted as “distance” in a Hotelling interpretation or “taste” in a random utility, discrete choice framework. We take a different approach to modeling (and varying) competition that allows us to hold preferences—and thus the potential social surplus—constant. We also reach different conclusions about the desirability of competition. For example, ? highlight a tradeoff from increasing competition when agents allocate effort between multiple, imperfectly observable or

⁸For recent contributions to this literature that assume perfectly competitive markets, see, e.g., ?, ?, and ?.

⁹? analyze a related labor market setting with adverse selection using an on-the-job search model, but their focus is quite different from ours.

¹⁰With private values, screening is useful only for rent extraction. Increasing competition reduces these rents, and hence the incentive to screen, causing welfare to rise. With common values, increasing competition strengthens incentives to separate, causing welfare to (eventually) decline.

¹¹Also see ?, ?, and ?, among others.

¹²Also see ?, who propose a slightly different model of imperfect competition.

contractible tasks. However, without multitasking, they find that competition improves welfare, even in the presence of adverse selection. This is also the case in ?, who restrict attention to single-price contracts. ? also restrict attention to a single contract, but with endogenous “quality,” and find that welfare is maximized under monopoly. In comparison to these papers, we find that competition can be beneficial or harmful. Though a number of differences (e.g., multidimensional heterogeneity, the contract space, the equilibrium concept) preclude a direct comparison, we interpret the results in these papers as providing a distinct but complementary insight about the interaction between competition and adverse selection.

2 Model

The Environment. We consider an economy with two buyers and a unit measure of sellers. Each seller is endowed with a single unit of a perfectly divisible good. Buyers have no capacity constraints, i.e., they can trade with many sellers. A fraction $\mu_l \in (0, 1)$ of sellers possess a low (l) quality good, while the remaining fraction $\mu_h = 1 - \mu_l$ possess a high (h) quality good. Buyers and sellers derive utility v_i and c_i , respectively, from consuming each unit of a quality $i \in \{l, h\}$ good, with $v_l < v_h$ and $c_l < c_h$. We assume that there are gains from trading both high- and low-quality goods, i.e., that $v_i > c_i$ for $i \in \{l, h\}$.

There are two types of frictions in the market. First, there is *asymmetric information*: sellers observe the quality of the good they possess while buyers do not, though the probability μ_i that a randomly selected good is quality $i \in \{l, h\}$ is common knowledge. In order to generate the standard “lemons problem,” we focus on the case in which $v_l < c_h$.

The second type of friction is a *search* friction: as we describe in detail below, the buyers in our model will make offers, but the sellers will not necessarily sample (or have access to) all offers. In particular, we assume that a fraction $1 - p$ of sellers will be matched with—and hence receive an offer from—a single buyer, which we assume is equally likely to be either buyer. The remaining fraction of sellers, p , will be matched with both buyers. A seller can only trade with a buyer if they are matched. Throughout the paper, we refer to sellers who are matched with one buyer as “captive,” since they only have one option for trade, and we refer to those who are matched with two buyers as “noncaptive.”

Given these search frictions, a buyer understands that, conditional on being matched with a particular

seller, this seller will be captive with probability $1 - \pi$ and noncaptive with probability π , where

$$\pi = \frac{p}{\frac{1}{2}(1-p) + p} = \frac{2p}{1+p}. \quad (1)$$

This formalization of search frictions is helpful for deriving and explaining our key results in the simplest possible manner. For one, it allows us to vary the degree of competition with a single parameter, π , nesting monoposony and perfect competition as special cases.¹³ Second, since the current formulation ensures that all sellers are matched with at least one buyer, a change in π varies the degree of competition without changing the potential gains from trade or “coverage” in the market; this is particularly helpful in isolating the effects of competition on welfare. However, it is important to stress that our equilibrium characterization and the ensuing results extend to markets with an arbitrary number of buyers and more general meeting technologies; see Section 7.

Offers, Payoffs, and Definition of Equilibrium. We model the interaction between a seller and the buyer(s) that she meets as a game in which the buyer(s) choose a mechanism and the seller chooses a message to send to each buyer she meets. A buyer’s mechanism is a function that maps the seller’s message into an offer, which specifies a quantity of numeraire to be exchanged for a certain fraction of the seller’s good.¹⁴ The seller’s message space can be arbitrarily large: it could include the quality of her good, whether or not she is in contact with the other buyer, the details of the other buyer’s mechanism, and any other (even not payoff-relevant) information. Importantly, we assume that mechanisms are exclusive, in the sense that a seller can choose to accept the offer generated by only one buyer’s mechanism, even when two offers are available.

In Appendix C, we apply insights from the delegation principle (??) to show that, in our environment, it is sufficient to restrict attention to menu games where buyers offer a menu of two contracts.¹⁵ In particular, letting x denote the quantity of good to be exchanged for t units of numeraire, a buyer’s offer can be summarized by the menu $\{(x_l, t_l), (x_h, t_h)\} \in \left([0, 1] \times \mathbb{R}_+\right)^2$, where (x_i, t_i) is the contract intended for a seller of type $i \in \{l, h\}$. A seller who owns a quality i good and accepts a contract (x, t) receives a payoff $t + (1 - x)c_i$, while a buyer who acquires a quality i good at terms (x, t) receives a

¹³Given the relationship in (1), it turns out that varying p or π is equivalent for all of our results below. We choose π because it simplifies some of the equations.

¹⁴The mechanisms we consider are assumed to be deterministic, but otherwise unrestricted. Stochastic mechanisms present considerable technical challenges and raise other conceptual issues that are, in our view, tangential to our key results.

¹⁵To be more precise, we show that the (distribution of) equilibrium allocations in any game where buyers offer the general mechanisms described above coincide with those in another game in which buyers only offer a menu of two contracts.

payoff $-t + xv_i$. Meanwhile, a seller with a quality i good who does not trade receives a payoff c_i , while a buyer who does not trade receives zero payoff.

Let $\mathbf{z}_i = (x_i, t_i)$ denote the contract that is intended for a seller of type $i \in \{l, h\}$, and let $\mathbf{z} = (\mathbf{z}_l, \mathbf{z}_h)$. A buyer's strategy, then, is a distribution across menus, $\Phi \in \Delta([0, 1] \times \mathbb{R}_+)^2$. A seller's strategy is much simpler: given the available menus, a seller chooses the menu with the contract that maximizes her payoffs or mixes between menus if she is indifferent. Conditional on a menu, the seller chooses the contract that maximizes her payoffs. In what follows, we will take the seller's optimal behavior as given.

A symmetric equilibrium is thus a distribution $\Phi^*(\mathbf{z})$ such that:

1. *Incentive compatibility*: for almost all $\mathbf{z} = \{(x_l, t_l), (x_h, t_h)\}$ in the support of $\Phi^*(\mathbf{z})$,

$$t_i + c_i(1 - x_i) \geq t_{-i} + c_i(1 - x_{-i}) \quad \text{for } i \in \{l, h\}. \quad (2)$$

2. *Buyer's optimize*: for almost all $\mathbf{z} = \{(x_l, t_l), (x_h, t_h)\}$ in the support of $\Phi^*(\mathbf{z})$,

$$\mathbf{z} \in \arg \max_{\mathbf{z}} \sum_{i \in \{l, h\}} \mu_i (v_i x_i - t_i) \left[1 - \pi + \pi \int_{\mathbf{z}'} \chi_i(\mathbf{z}, \mathbf{z}') \Phi^*(d\mathbf{z}') \right], \quad (3)$$

where

$$\chi_i(\mathbf{z}, \mathbf{z}') = \begin{cases} 0 \\ \frac{1}{2} \\ 1 \end{cases} \quad \text{if} \quad t_i + c_i(1 - x_i) \begin{cases} < \\ = \\ > \end{cases} t'_i + c_i(1 - x'_i). \quad (4)$$

The function χ_i reflects the seller's optimal choice. We assume that a seller who is indifferent between menus randomizes with equal probability. Within a given menu, we assume that sellers do not randomize; for any incentive compatible contract, sellers choose the contract intended for their type, as in most of the mechanism design literature (see, e.g., ?, ?).

3 Properties of Equilibria

Characterizing the equilibrium described above requires solving for a distribution over four-dimensional menus. In this section, we first show that each menu can be summarized by the indirect utilities that it delivers to each type of seller, so that equilibrium strategies can be defined by a joint distribution over two-dimensional objects. Then, we establish that the marginal distributions of offers intended for each type of seller have fully connected support and no mass points. Finally, we establish that, in equilibrium,

the two contracts offered by any buyer are ranked in exactly the same way by both low- and high-type sellers: a low-type seller prefers buyer 1's offer to buyer 2's offer if and only if a high-type seller does as well. This property of equilibria, which we call "strictly rank-preserving," simplifies the characterization even more, as the marginal distribution of offers for high-quality sellers can be expressed as a strictly monotonic transformation of the marginal distribution of offers for low-quality sellers.

3.1 Utility Representation

As a first step, we establish that any menu can be summarized by two numbers, (u_l, u_h) , where $u_i = t_i + c_i(1 - x_i)$ denotes the utility received by a type $i \in \{l, h\}$ seller from accepting a contract z_i .

Lemma 1. *In any equilibrium, for almost all $z \in \text{supp}(\Phi^*)$, it must be that $x_l = 1$ and $t_l = t_h + c_l(1 - x_h)$.*

In words, Lemma 1 states that all equilibrium menus require that low-quality sellers trade their entire endowment, and that their incentive compatibility constraint always binds. This is reminiscent of the "no-distortion-at-the-top" result in the taxation literature, or that of full insurance for the high-risk agents in ?.

Corollary 1. *In equilibrium, any menu of contracts $\{(x_l, t_l), (x_h, t_h)\} \in ([0, 1] \times \mathbb{R}_+)^2$ can be summarized by a pair (u_l, u_h) with $x_l = 1$, $t_l = u_l$,*

$$x_h = 1 - \frac{u_h - u_l}{c_h - c_l}, \text{ and} \quad (5)$$

$$t_h = \frac{u_l c_h - u_h c_l}{c_h - c_l}. \quad (6)$$

Notice that, since $0 \leq x_h \leq 1$, feasibility requires that the pair (u_l, u_h) satisfies

$$c_h - c_l \geq u_h - u_l \geq 0. \quad (7)$$

In what follows, we will often refer to the requirement $u_h \geq u_l$ as a "monotonicity constraint." Note that, when this constraint binds, Corollary 1 implies that $x_h = 1$ and $t_h = t_l$.

3.2 Recasting the Buyer's Problem and Equilibrium

Lemma 1 and Corollary 1 allow us to recast the problem of a buyer as choosing a menu of indirect utilities, (u_l, u_h) , taking as given the distribution of indirect utilities offered by the other buyer. For any

menu (u_l, u_h) , a buyer must infer the probability that the menu will be accepted by a type $i \in \{l, h\}$ seller. In order to calculate these probabilities, let us define the marginal distributions

$$F_i(u_i) = \int_{z'_i} \mathbf{1} [t'_i + c_i (1 - x'_i) \leq u_i] \Phi(dz'_i)$$

for $i \in \{l, h\}$. In words, $F_l(u_l)$ and $F_h(u_h)$ are the probability distributions of indirect utilities arising from each buyer's mixed strategy. When these distributions are continuous and have no mass points, the probability that a contract intended for a type i seller is accepted is simply $1 - \pi + \pi F_i(u_i)$, i.e., the probability that the seller is captive plus the probability that he is noncaptive but receives another offer less than u_i . However, if $F_i(\cdot)$ has a mass point at u_i , then the fraction of noncaptive sellers of type i attracted to a contract with value u_i is given by $\tilde{F}_i(u_i) = \frac{1}{2}F_i^-(u_i) + \frac{1}{2}F_i(u_i)$, where $F_i^-(u_i) = \lim_{u \nearrow u_i} F_i(u)$ is the left limit of F_i at u_i . Given $\tilde{F}_i(\cdot)$, a buyer solves

$$\max_{u_l \geq c_l, u_h \geq c_h} \mu_l (1 - \pi + \pi \tilde{F}_l(u_l)) \Pi_l(u_l, u_h) + \mu_h (1 - \pi + \pi \tilde{F}_h(u_h)) \Pi_h(u_l, u_h) \quad (8)$$

subject to (7), with

$$\Pi_l(u_l, u_h) \equiv v_l x_l - t_l = v_l - u_l \quad (9)$$

$$\Pi_h(u_l, u_h) \equiv v_h x_h - t_h = v_h - u_h \frac{v_h - c_l}{c_h - c_l} + u_l \frac{v_h - c_h}{c_h - c_l}. \quad (10)$$

In words, $\Pi_i(u_l, u_h)$ is the buyer's payoff conditional on the offer u_i being accepted by a type i seller. We refer to the objective in (8) as $\Pi(u_l, u_h)$.

Before proceeding, note that $\Pi_h(u_l, u_h)$ is increasing in u_l : by offering more utility to low-quality sellers, the buyer relaxes the incentive constraint and can earn more profits when he trades with high-quality sellers. As a result, one can easily show that the profit function $\Pi(u_l, u_h)$ is (at least) weakly supermodular. This property will be important in several of the results we establish below.

Using the optimization problem described above, we can redefine the equilibrium in terms of the distributions of indirect utilities. In particular, for each u_l , let

$$U_h(u_l) = \{u'_h \in \arg \max \Pi(u_l, u'_h) \mid u'_h \geq c_h \cap c_h - c_l \geq u'_h - u_l \geq 0\}.$$

The equilibrium can then be described by the marginal distributions $\{F_i(u_i)\}_{i \in \{l, h\}}$ together with the requirement that a joint distribution function must exist. In other words, a probability measure Φ over

the set of feasible pairs (u_l, u_h) must exist such that, for each $u_l > u'_l$ and $u_h > u'_h$

$$1 = \Phi(\{(\hat{u}_l, \hat{u}_h); \hat{u}_h \in U_h(\hat{u}_l), \hat{u}_l \in [c_l, v_h]\})$$

$$F_l^-(u_l) - F_l(u'_l) = \Phi(\{(\hat{u}_l, \hat{u}_h); \hat{u}_h \in U_h(\hat{u}_l), \hat{u}_l \in (u'_l, u_l)\}), \quad (11)$$

$$F_h^-(u_h) - F_h(u'_h) = \Phi(\{(\hat{u}_l, \hat{u}_h); \hat{u}_h \in U_h(\hat{u}_l), \hat{u}_h \in (u'_h, u_h)\}). \quad (12)$$

Note that this definition imposes two requirements. The first is that buyers behave optimally: for each u_l , the joint probability measure puts a positive weight only on $u_h \in U_h(u_l)$. The second is aggregate consistency, i.e., that F_l and F_h are marginal distributions associated with a joint measure of menus.

3.3 Basic Properties of Equilibrium Distributions

In this section, we establish that, in equilibrium, the distributions $F_l(u_l)$ and $F_h(u_h)$ are continuous and have connected support, i.e., there are neither mass points nor gaps in either distribution.

Proposition 1. *The marginal distributions F_l and F_h have connected support. They are also continuous, with the possible exception of a mass point in F_l at v_l .*

As in ?, the proof of Proposition 1 rules out gaps and mass points in the distribution by constructing profitable deviations. A complication that arises in our model, which does not arise in ?, is that payoffs are interdependent, e.g., a change in the utility offered to low-quality sellers changes the contract—and hence the profits—that a buyer receives from high-quality sellers. We prove the properties of F_l and F_h described in Proposition 1 sequentially: we first show that F_h is continuous and strictly increasing, and then apply an inductive argument to prove that F_l has connected support and is continuous, with a possible exception at the lower bound of the support. An important step in the induction argument, which we later use more generally, is to show that the objective function $\Pi(u_l, u_h)$ is strictly supermodular. We state this here as a lemma.

Lemma 2. *The profit function is strictly supermodular, i.e.,*

$$\Pi(u_{i1}, u_{h1}) + \Pi(u_{i2}, u_{h2}) \geq \Pi(u_{i2}, u_{h1}) + \Pi(u_{i1}, u_{h2}), \quad \forall u_{i1} \geq u_{i2}, \quad i \in \{l, h\}$$

with strict inequality when $u_{i1} > u_{i2}$, $i \in \{l, h\}$.

As noted above, the supermodularity of the buyer's profit function reflects a basic complementarity between the indirect utilities offered to low- and high-quality sellers. An important implication is that

the correspondence $U_h(u_l)$ is weakly increasing. We use this property to construct deviations to rule out gaps and mass points in the distribution F_l almost everywhere in its support; later, in Section 4, we show that these mass points only occur in a knife-edge case. Hence, generically, the marginal distribution F_l has connected support and no mass points in its support.

3.4 Strict rank-preserving

In this section, we establish that every equilibrium has the property that the menus being offered are *strictly rank-preserving*—that is, low- and high-quality sellers share the same ranking over the set of menus offered in equilibrium—with the possible exception of the knife-edge case discussed above. We prove this result by showing that the mapping between a buyer’s optimal offer to low- and high-quality sellers, $U_h(u_l)$, is a well defined, strictly increasing function. We start with the following definition.

Definition 1. For any subset U_l of $\text{Supp}(F_l)$, an equilibrium is *strictly rank-preserving over U_l* if the correspondence $U_h(u_l)$ is a strictly increasing function of u_l for all $u_l \in U_l$. An equilibrium is *strictly rank-preserving* if it is strictly rank-preserving over $\text{Supp}(F_l)$.

Equivalently, an equilibrium is strictly rank-preserving if, for any (u_l, u_h) and (u'_l, u'_h) in the equilibrium support, $u_l > u'_l$ if and only if $u_h > u'_h$. Given this terminology, we now establish a key result.

Theorem 1. All equilibria are strictly rank-preserving over the set $\text{Supp}(F_l) \setminus \{v_l\}$.

Theorem 1 follows from the facts established above. In particular, the strict supermodularity of $\Pi(u_l, u_h)$ implies that $U_h(u_l)$ is a weakly increasing correspondence. However, since $F_l(\cdot)$ and $F_h(\cdot)$ are strictly increasing and continuous, we show that $U_h(u_l)$ can neither be multi-valued nor have flats. Intuitively, if there exists a $u_l > \min \text{Supp}(F_l) \equiv \underline{u}_l$ and $u'_h > u_h$ such that $u_h, u'_h \in U_h(u_l)$, then the supermodularity of $\Pi(u_l, u_h)$ implies that $[u_h, u'_h] \subset U_h(u_l)$. Since $F_h(\cdot)$ has connected support, if U_h were a correspondence for some u_l , then this would imply that $F_l(\cdot)$ must have a mass point at u_l , which contradicts Proposition 1. Similarly, if there exists u_h and $u'_l > u_l$ offered in equilibrium such that $U_h(u'_l) = U_h(u_l) = u_h$, then F_h would feature a mass point, in contradiction with Proposition 1. Hence, $U_h(u_l)$ must be a strictly increasing *function* for all $u_l > \underline{u}_l$.

If $F_l(\cdot)$ is continuous everywhere, then every menu offered in equilibrium is accepted by the same fraction of low- and high-quality noncaptive sellers; we state this formally in the following Corollary.

Corollary 2. If F_l and F_h are continuous, then $F_h(U_h(u_l)) = F_l(u_l)$.

Taken together, Theorem 1 and Corollary 2 simplify the construction of an equilibrium. Specifically, when an equilibrium exists in which the marginal distributions are continuous, then the equilibrium can be described compactly by the marginal distribution $F_l(u_l)$ and the function $U_h(u_l)$.

4 Construction of Equilibrium

In this section, we use the properties established above to construct equilibria. Then, we show that the equilibrium we construct is unique. In this sense, we characterize the entire set of equilibrium outcomes.

4.1 Special Cases: Monopsony and Perfect Competition

To fix ideas, we first characterize equilibria in the well-known special cases of $\pi = 0$ and $\pi = 1$, i.e., when sellers face a monopsonist and when they face two buyers in Bertrand competition, respectively. As we will see, several features of the equilibrium in these two extreme cases guide our construction of equilibria for the general case of $\pi \in (0, 1)$.

When $\pi = 0$, so that each seller meets with at most one buyer, the buyers solve

$$\max_{(u_l, u_h)} \mu_l(v_l - u_l) + \mu_h \left[v_h - u_h \frac{v_h - c_l}{c_h - c_l} + u_l \frac{v_h - c_h}{c_h - c_l} \right],$$

subject to the monotonicity and feasibility constraints in (7). The solution to this problem, summarized in Lemma 3 below, is standard and, hence, we omit the proof.

Lemma 3. *Suppose $\pi = 0$, and let*

$$\phi_l \equiv 1 - \frac{\mu_h}{\mu_l} \left(\frac{v_h - c_h}{c_h - c_l} \right). \quad (13)$$

If $\phi_l > 0$, then the unique equilibrium has $u_l = c_l$ with $x_l = 1$ and $u_h = c_h$ with $x_h = 0$; if $\phi_l < 0$, then $u_l = u_h = c_h$ with $x_l = x_h = 1$; and if $\phi_l = 0$, then $u_l \in [c_l, c_h]$ with $x_l = 1$ and $u_h = c_h$ with $x_h \in [0, 1]$.

The parameter ϕ_l is a summary statistic for the adverse selection problem: it represents the net marginal cost (to the buyer) of delivering an additional unit of utility to a low-quality seller. It is strictly less than 1 because the direct cost of an additional unit of transfer to a low-quality seller is partially offset by the indirect benefit of relaxing this seller's incentive constraint, which allows the buyer to trade more with a high-quality seller. This indirect benefit is captured by the second term on the right-hand side: when this term is large, ϕ_l is small, the cost of trading with high-quality sellers is low, and adverse selection is *mild*. Conversely, when this term is small, ϕ_l is large, it is costly to trade with high-quality

sellers, and therefore adverse selection is relatively *severe*. According to this measure, adverse selection is thus severe when the relative fraction of high-quality sellers, μ_h/μ_l , is small; the gains from trading with high-quality sellers, $v_h - c_h$, are relatively small; and/or the information rents associated with separating high- and low-quality sellers, $c_h - c_l$, are large.

When $\phi_l > 0$, the net cost to a buyer of increasing u_l is positive, so she sets u_l as low as possible, i.e., $u_l = c_l$. This implies that the high-quality seller is entirely shut out, i.e., $x_h = 0$. Otherwise, when $\phi_l < 0$, increasing u_l yields a net benefit to the buyer. As a result, a buyer raises u_l until the monotonicity constraint in (7) binds, i.e., she pools high- and low-quality sellers, offering $u_h = u_l = c_h$.

Before proceeding to the perfectly competitive case, we highlight two features of the equilibrium under monopsony. First, buyers offer *separating menus* ($u_h > u_l$) when ϕ_l is positive and *pooling menus* ($u_h = u_l$) when ϕ_l is negative. Second, they make non-negative payoffs on both types when $\phi_l > 0$, but lose money on low-quality sellers when $\phi_l < 0$. In other words, the equilibrium features *cross-subsidization* when ϕ_l is negative, but not when ϕ_l is positive.

When competition is perfect, i.e., when $\pi = 1$, our setup becomes the same as that in ?, and similar to that of ?. In this case, when $\phi_l \geq 0$, the unique equilibrium is in pure strategies, with buyers offering the standard “least-cost separating” contract; type l sellers earn $u_l = v_l$ and type h sellers trade a fraction of their endowment at a unit price of v_h , such that the incentive constraint of the low-quality seller binds. However, when $\phi_l < 0$, there is no pure strategy equilibrium.¹⁶ In this case, an equilibrium in mixed strategies emerges, as in ? and ?.¹⁷ Each buyer mixes over menus, all of which involve negative profits from low-quality sellers, offset exactly by positive profits from high-quality sellers, leading to zero profits. The marginal distribution $F_l(\cdot)$ is such that profitable deviations are ruled out. The following lemma summarizes these results.

Lemma 4. *When $\pi = 1$, the unique equilibrium is: (i) if $\phi_l \geq 0$, then $u_l = v_l$, $x_l = 1$, $u_h = \frac{v_h(c_h - c_l) + v_l(v_h - c_h)}{v_h - c_l}$, and $x_h = \frac{v_l - c_l}{v_h - c_l}$; (ii) if $\phi_l < 0$, then the symmetric equilibrium is described by the distribution*

$$F_l(u_l) = \left(\frac{u_l - v_l}{\mu_h(v_h - v_l)} \right)^{-\phi_l}, \quad (14)$$

¹⁶All buyers offering the least-cost separating contract cannot be an equilibrium, as a pooling offer is profitable. All buyers offering a pooling contract cannot be an equilibrium either, since it is vulnerable to a cream-skimming deviation, wherein a competing buyer draws away high-quality sellers by offering a contract with $x < 1$ but at a higher price.

¹⁷? shows that the equilibrium is unique.

with $\text{Supp}(F_l) = [v_l, \bar{v}]$ and $F_h(u_h) = F_l(U_h(u_l))$, where $\bar{v} = \mu_h v_h + \mu_l v_l$ and $U_h(u_l)$ satisfies

$$\mu_h \Pi_h(u_l, U_h(u_l)) + \mu_l \Pi_l(u_l, U_h(u_l)) = 0. \quad (15)$$

As with $\pi = 0$, equilibrium when $\pi = 1$ features no cross-subsidization when $\phi_l \geq 0$ and cross-subsidization when $\phi_l < 0$. However, unlike the case with $\pi = 0$, equilibrium with $\pi = 1$ features separating contracts for all values of ϕ_l . These properties guide our construction of equilibria in the next section, when we study the general case of $\pi \in (0, 1)$.

4.2 General Case: Imperfect Competition

We now describe how to construct equilibria when $\pi \in (0, 1)$. Recall that an equilibrium is summarized by a distribution $F_l(u_l)$ and a strictly increasing function $U_h(u_l)$. A key determinant of the structure of equilibrium menus is whether the monotonicity constraint in (7) is binding. When it is slack, the local optimality (or first-order) condition for u_l , along with the strict rank-preserving condition that relates $F_h(U_h(u_l)) = F_l(u_l)$ together characterize the equilibrium distribution $F_l(u_l)$. The function $U_h(u_l)$ then follows from the requirement that all menus $(u_l, U_h(u_l))$ must yield the buyer equal profits. When the monotonicity constraint is binding, the policy function is, by definition, $U_h(u_l) = u_l$.¹⁸

Our analysis of $\pi = 0$ and $\pi = 1$ points to the importance of ϕ_l . Recall that when $\phi_l > 0$, the monotonicity constraint was always slack, i.e., $U_h(u_l) > u_l$ for all $u_l \in \text{Supp}(F_l)$. When $\phi_l < 0$, on the other hand, the monotonicity constraint was binding only when $\pi = 0$ and slack at $\pi = 1$. Guided by these results, we discuss our construction separately for the $\phi_l > 0$ and the $\phi_l < 0$ cases.¹⁹

Case 1: $\phi_l > 0$. Given the analysis of $\pi = 0$ and $\pi = 1$, we conjecture that the monotonicity constraint is slack for any $\pi \in (0, 1)$. Proposition 2 confirms that this is indeed the case.

Proposition 2. *For any $\pi \in (0, 1)$ and $\phi_l > 0$, there exists an equilibrium in which (i) F_l satisfies the differential equation*

$$\frac{\pi f_l(u_l)}{1 - \pi + \pi F_l(u_l)} (v_l - u_l) = \phi_l, \quad (16)$$

¹⁸Of course, $u_h = U_h(u_l)$ must be locally optimal as well, but this condition is implied by the joint requirements on u_l and $U_h(u_l)$ described above.

¹⁹The equilibrium when $\phi_l = 0$ has a slightly different structure and, for the sake of brevity, we relegate analysis of this knife-edge case to Appendix D.

with the boundary condition $F_l(c_l) = 0$; and (ii) $U_h(u_l) > u_l$ and satisfies the equal profit condition

$$(1 - \pi + \pi F_l(u_l)) [\mu_h \Pi_h(u_l, U_h(u_l)) + \mu_l (v_l - u_l)] = \mu_l (1 - \pi) (v_l - c_l). \quad (17)$$

Equation (16) is derived by taking the first-order condition of (8) with respect to u_l —holding u_h fixed—and then imposing the strict rank-preserving property.²⁰ This necessary condition is familiar from basic production theory. The left-hand side is the marginal benefit to the buyer of increasing u_l , i.e., the product of the semi-elasticity of demand and the profit per trade. The right-hand side, ϕ_l , is the marginal cost of increasing the utility of the low-quality seller, taking into account the fact that increasing u_l relaxes the incentive constraint.²¹ Note that, even though (16) ensures that local deviations by a buyer from an equilibrium menu are not profitable, completing the proof requires ensuring that there are no profitable global deviations as well; we establish that this is true in Appendix A.2.1.

The boundary condition requires that the lowest utility offered to the low-quality seller is c_l . From (17), and the fact that $F_l(c_l) = 0$, we find $U_h(c_l) = c_h$, so that the worst menu offered in equilibrium coincides with the monopsony outcome. Intuitively, if the worst menu offers more utility to low-quality sellers than c_l , the buyer could profit by decreasing u_l and u_h ; the gains associated with trading at better terms with the low types would exceed the losses associated from trading less quantity with high types, precisely because $\phi_l > 0$. Given that $\underline{u}_l = c_l$, if the worst equilibrium menu offers more utility to high-quality sellers than c_h , then a buyer offering this menu could profit by decreasing u_h ; his payoff from trading with high types would increase without changing the payoffs from trading with low types.

The final equilibrium object, $U_h(u_l)$, is characterized by the equal profit condition: the left side of (17) defines the buyer's payoff from the menu $(u_l, U_h(u_l))$, while the right side is the profit earned from the worst contract offered in equilibrium. Figure 1 plots the two equilibrium functions in this region.

Notice from (16) that, since $\phi_l > 0$, our equilibrium has $v_l > u_l$ for all menus in equilibrium, so that buyers earn strictly positive profits from trading with low-quality sellers. It is straightforward to show that buyers also earn strictly positive profits from trading with high-quality sellers. Hence, in this region, the equilibrium features no cross-subsidization, as was the case for $\pi = 0$ and $\pi = 1$. Finally, it is also worth noting that the equilibrium distribution of offers converges to the limiting cases as π converges to both 0 and 1; in the former case, the distribution converges to a mass point at the monopsony outcome,

²⁰As we discuss in the proof of Proposition 2, this first-order condition requires three assumptions: that $u_h > u_l$ for all menus; that there is no mass point at the lower bound of the support of $F_l(u_l)$; and that the implied quantity traded by the high-quality seller is interior in all trades, i.e., $0 < x_h = (u_h - u_l)/(c_h - c_l) < 1$, except possibly at the boundary of the support of F_l . All of these assumptions are confirmed in equilibrium.

²¹It is straightforward to derive a closed-form solution for $F_l(u_l)$ from (16); see equation (1) in the Appendix.

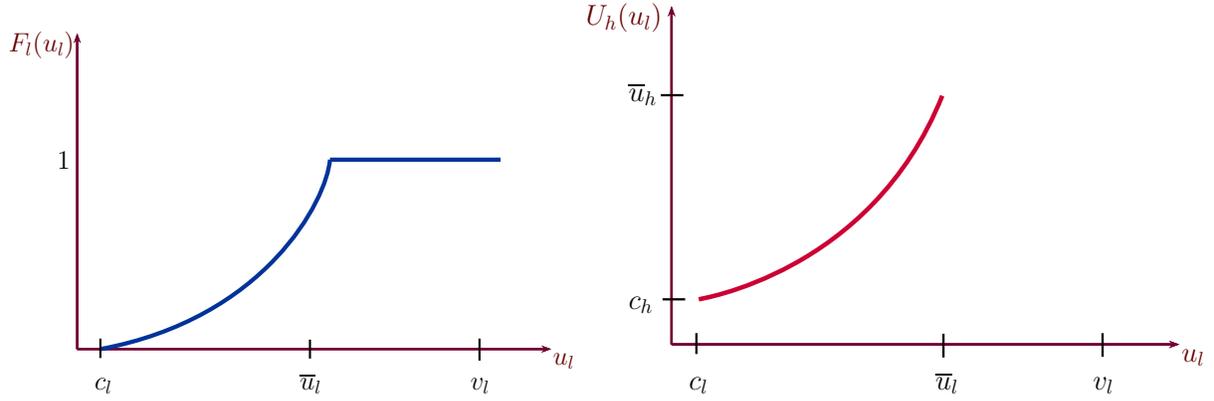


Figure 1: Equilibrium for $\pi \in (0, 1)$, $\phi_l > 0$. The left panel plots the CDF $F_l(u_l)$ and the right panel plots the mapping $U_h(u_l)$.

while in the latter case, the distribution converges to a mass point at the least-cost separating outcome.

Case 2: $\phi_l < 0$. In this region of the parameter space, the equilibrium features a pooling menu when $\pi = 0$ and a distribution of separating menus when $\pi = 1$. This leads us to conjecture that the equilibrium for $\pi \in (0, 1)$ can feature pooling, separating, or a mixture of the two, depending on the value of π . The following lemma formalizes this conjecture and shows the existence of a threshold utility for the offer to low-quality sellers, \hat{u}_l , such that all offers with $u_l < \hat{u}_l$ are pooling menus, while all offers with $u_l > \hat{u}_l$ are separating menus.²² Depending on whether this threshold lies at the lower bound, the upper bound, or in the interior of the support of $F_l(u_l)$, there are three possible cases, respectively: all equilibrium offers are separating menus, all are pooling menus, or there is a mixture with some pooling menus (offering relatively low utility to the seller) and some separating menus (offering higher utility). Later, in Proposition 4, we provide conditions on ϕ_l and π under which each case obtains.

Proposition 3. *For any $\pi \in (0, 1)$ and $\phi_l < 0$, there exists an equilibrium in which: (i) there exists a threshold \hat{u}_l such that, for any u_l in the interior of $\text{Supp}(F_l)$:*

1. if $u_l \leq \hat{u}_l$, $U_h(u_l) = u_l$ and F_l satisfies

$$\frac{\pi f_l(u_l)}{1 - \pi + \pi F_l(u_l)} (\mu_h v_h + \mu_l v_l - u_l) = 1, \quad (18)$$

2. if $u_l > \hat{u}_l$, $U_h(u_l) > u_l$ and F_l satisfies (16).

²²At this point, it may seem arbitrary to conjecture that pooling occurs at the bottom of the distribution and separation at the top. As we will discuss later in the text, the reason this is ultimately true is that the cream-skimming deviation—which makes the pooling offer suboptimal—becomes more attractive as the indirect utility being offered increases.

(ii) $U_h(\underline{u}_l) = c_h$ and $U_h(\bar{u}_l) = \bar{u}_l$.

To understand the first set of (necessary) conditions in Proposition 3, consider the region where the buyers offer pooling menus. Here, buyers trade off profit per trade against the probability of trade, with no interaction between offers and incentive constraints. As a result, the equilibrium in this pooling region behaves as in the canonical single-quality model, with the buyer's payoff equal to the average value $\mu_h v_h + (1 - \mu_h)v_l$. This yields (18). In the region where buyers offer separating menus, $F_l(u_l)$ is characterized by the local optimality condition (16), exactly as in the $\phi_l > 0$ case. Recall from our discussion that this differential equation accounts explicitly for the effect of an offer u_l on the seller's incentive constraint. In this region, $U_h(u_l)$ is determined by the equal profit condition.

The second part of the result describes boundary conditions for the worst and best menus offered in equilibrium. The first condition requires that the worst menu yields utility c_h to high-quality sellers. To see why, suppose the worst menu is a pooling menu with $U_h(\underline{u}_l) = \underline{u}_l > c_h$. Then, lowering both u_h and u_l leads to strictly higher profits. If the worst menu is separating with $U_h(\underline{u}_l) > c_h$, then a downward deviation in only u_h is feasible and strictly increases profits. The second condition requires that the best menu offered in equilibrium is a pooling menu. Intuitively, if the best menu offered in equilibrium were a separating menu, then $\chi_h < 1$. This cannot be optimal when $\phi_l < 0$: the buyer can trade more with the high-quality seller by increasing the utility offered to low-quality sellers. Since this is already the best menu in equilibrium, this deviation has no impact on the number of sellers the buyer attracts but yields strictly higher profits.

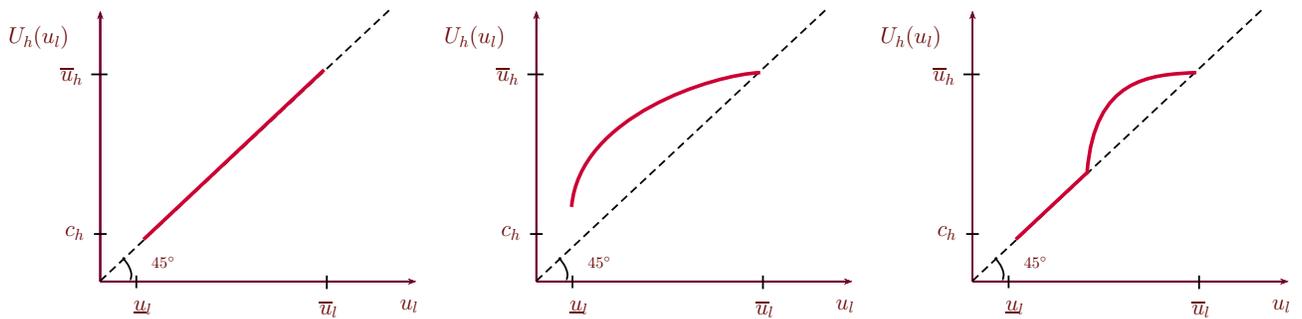


Figure 2: The mapping $U_h(u_l)$ for all pooling, all separating, and mixed equilibria when $\phi_l < 0$

Given these properties, we now establish two critical values— $\phi_1(\pi)$ and $\phi_2(\pi)$, with $\phi_2(\pi) < \phi_1(\pi) < 0$ —that determine which of the three cases described above emerge in equilibrium. When $\phi_l < \phi_2(\pi)$, the threshold $\hat{u}_l = \bar{u}_l$ and there is an *all pooling* equilibrium. When $\phi_l > \phi_1(\pi)$, the monotonicity

constraint is slack almost everywhere, so that $\hat{u}_l = \underline{u}_l$, and the equilibrium features *all separating* menus. Finally, if ϕ_l lies between these two critical values, we have a *mixed* equilibrium, with an intermediate threshold $\hat{u}_l \in (\underline{u}_l, \bar{u}_l)$. Figure 2 illustrates $\mathcal{U}_h(u_l)$ for all three possibilities.

Proposition 4. *For any $\pi \in (0, 1)$, there exist cutoffs $\phi_2(\pi) < \phi_1(\pi) < 0$ such that an all pooling equilibrium exists for all $\phi_l \leq \phi_2(\pi)$, a mixed equilibrium exists for all $\phi_l \in (\phi_2(\pi), \phi_1(\pi))$, and an all separating equilibrium exists for all $\phi_l \in (\phi_1(\pi), 0)$.*

Intuitively, for a pooling menu (u_l, u_l) to be offered in equilibrium, the cream-skimming deviation $(u_l - \varepsilon, u_l)$ for some $\varepsilon > 0$ cannot yield strictly higher profits. To see how incentives to cream-skim vary with ϕ_l and π , notice that there are two sources of higher profits from the menu $(u_l - \varepsilon, u_l)$, relative to the candidate pooling menu. First, it decreases the loss conditional on trading with a low-quality seller. Second, it reduces the probability of trading with a noncaptive low-quality seller; since the buyer loses money on these sellers, this reduction in trading probability raises profits. The cost of cream-skimming is that the buyer earns lower profits on high-quality sellers. Therefore, incentives to cream-skim are weak—and thus pooling is easier to sustain—when high-quality sellers are relatively abundant (ϕ_l very negative) and/or there are relatively few noncaptive sellers (π is small).

The higher the level of utility being offered in a pooling menu, the more vulnerable it is to cream-skimming. Hence, if such a deviation is profitable at the lowest candidate value, c_h , then pooling cannot be sustained at all: this is the condition that determines the cutoff $\phi_1(\pi)$. Similarly, the cutoff $\phi_2(\pi)$ defines the boundary at which cream-skimming is not profitable even at the best pooling menu, \bar{u}_l . We derive these thresholds formally and provide a full equilibrium characterization in Appendix A.2.2.

Notice that, in all three cases, $u_l > v_l$ (since $u_l \geq c_h > v_l$) so that buyers always suffer losses when trading with low-quality sellers. Hence, as in the extreme cases of $\pi = 0$ and $\pi = 1$, there is cross-subsidization in every equilibrium when $\phi_l < 0$. Finally, as in the case of $\phi_l > 0$, the equilibrium distribution converges to the limiting cases as π converges to both 0 and 1.

Figure 3 summarizes the various types of equilibria and the regions in which each one obtains. The x- and y-axes represent the intensity of competition and severity of adverse selection, respectively. Recall that the latter is summarized by ϕ_l , which is a function of μ_h , the fraction of high-quality goods, as well as the valuations v_h, c_h, c_l . For concreteness, we use μ_h to vary ϕ_l on the y-axis—a higher fraction of high-quality goods implies a lower ϕ_l and, therefore, milder adverse selection.²³

²³The boundaries are redefined accordingly: $\mu_h \leq \mu_0$ if and only if $\phi_l \geq 0$ and $\mu_h \leq \mu_j(\pi)$ if and only if $\phi_l \geq \phi_j(\pi)$, $j \in \{1, 2\}$.

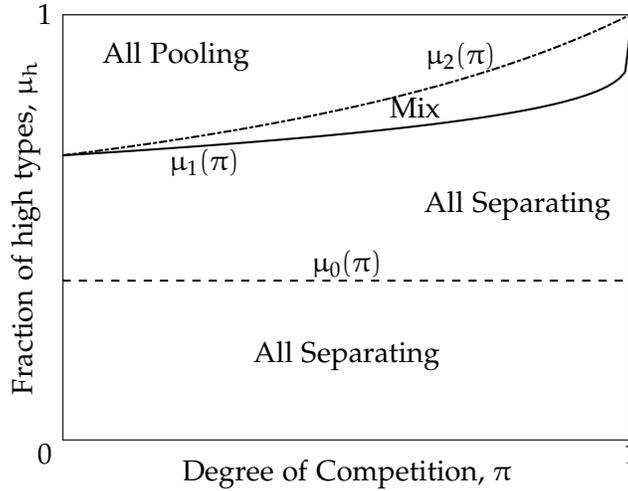


Figure 3: Equilibrium regions

4.3 Uniqueness

In Theorem 2, below, we establish that the equilibria constructed above are unique. For intuition, we sketch the arguments here for $\phi_l \neq 0$.²⁴ First, we show that equilibria do not feature a mass point, even at v_l . Next, when $\phi_l > 0$, we prove that equilibria do not feature pooling menus on a positive measure subset of F_l ; since equilibria have no mass points and must be separating almost everywhere, the equilibrium we construct in Proposition 2 describes the unique equilibrium. Finally, when $\phi_l < 0$, we demonstrate uniqueness in steps. First, we show that any equilibrium features pooling at the upper bound of the support of F_l . Second, we prove that any equilibrium features at most one interval of pooling menus followed by at most one interval of separating menus. Third, we prove that the equilibria characterized in Proposition 4 are mutually exclusive, so that equilibria without mass points are unique. Since no equilibrium features mass points when $\phi_l < 0$, these results establish the uniqueness of the equilibrium characterized in Proposition 4.

Theorem 2. *For any $\pi \in (0, 1)$ and $\phi_l \in \mathbb{R}$, there exists an equilibrium and it is unique.*

Note that we obtain a unique equilibrium without any refinements or restrictions on off-path behavior. This is because buyers' payoffs are well defined for any offer. In particular, since buyers are not capacity-constrained, the fraction of type $i \in \{l, h\}$ sellers who accept any offer (u_l, u_h) is uniquely determined by the (exogenous) meeting technology and the (endogenous) distribution of offers.²⁵

²⁴In Appendix D, we also prove uniqueness for the knife-edge case of $\phi_l = 0$.

²⁵Refinements are often necessary in models with capacity-constrained buyers, when two types of sellers would like to accept an off-path offer, and the probability that each type is able to execute the trade is not pinned down.

4.4 Discussion

The equilibrium characterized above has a number of testable implications. To start, we highlight three robust predictions about the properties of equilibrium menus. First, the strict rank-preserving property suggests a positive correlation between the contracts that buyers offer to different types of sellers: those buyers who make attractive offers to low-quality sellers will also make attractive offers to high-quality sellers. Hence, in equilibrium, buyers do not specialize in trading with a particular type of seller, but rather trade with equal frequency across all types. Second, whether buyers pool different types of sellers or separate them (using a menu of options) depends crucially on the severity of the two frictions.²⁶ Separation is more likely when adverse selection is relatively severe—so that the information costs of trading with high-quality sellers are large relative to the benefits—and competition is relatively strong—so that the payoffs from cream-skimming are relatively high.²⁷ Pooling is more likely when competition among buyers is weak and adverse selection is mild. Third, the theory predicts that menus that are less attractive, from the perspective of sellers, are more likely to be pooling. In other words, those who are posting offers with relatively unattractive terms should be offering fewer options and should account for a smaller share of observed transactions.

Our analysis also has implications about dispersion. In the region with separating menus, the model predicts dispersion *within* and *across* types. This is true both for quantities traded (coverage in an insurance context or loan size in a credit market context) as well as prices (premiums or interest rates, respectively). The extent of dispersion—both the support and the standard deviation of the quantity/price distributions—is determined by the interaction of competition (measured by π) and adverse selection (measured by ϕ_1). This joint dependence calls into question the practice of identifying imperfect competition or asymmetric information in isolation using cross-sectional dispersion. For example, a common empirical strategy to identify adverse selection is to test the correlation between the quantity an agent trades and her type, as measured by ex-post outcomes.²⁸ In our equilibrium, depending on market structure, the correlation between the seller's quality and the quantity she sells can be either negative or zero. As a result, using the relationship between quantity and type without accounting for the imperfect nature of competition is likely to yield misleading conclusions. A similar concern applies to the strategy

²⁶This result stands in stark contrast to, e.g., ?. In that model, and many like it, the quantity traded with high-quality sellers is *independent* of the distribution of types in the market; trade with high-quality sellers is distorted even if the fraction of low-quality goods in the market is arbitrarily small.

²⁷Consistent with our findings, ? find evidence of greater cream-skimming by health insurance providers when the market is more competitive.

²⁸This technique for identifying adverse selection has been applied to a number of markets, following the seminal paper by ?; recent examples include ?, ?, and ?.

of identifying search frictions from price dispersion.²⁹ In markets where adverse selection is a concern, the magnitude of cross-sectional variation in terms of trade is also a function of selection-related parameters. Obtaining an accurate assessment of trading frictions in such settings thus requires controlling for the underlying distribution of types.

5 Increasing Competition and Reducing Information Asymmetries

Many markets in which adverse selection is a first-order concern are experiencing dramatic changes. Some of these changes are regulatory in nature; for example, as we describe in greater detail below, there are several recent policy initiatives to make health insurance markets and over-the-counter markets for financial securities more competitive and transparent. Other changes derive from technological improvements; for example, advances in credit scoring reduce information asymmetries in loan markets.

In this section, we use the framework developed above to examine the likely effects of these types of changes on economic activity. Our metric for economic activity is the utilitarian welfare function, which measures the expected gains from trade that are realized in equilibrium. We show that increasing competition or reducing information asymmetries can worsen the distortions from adverse selection—thereby *decreasing* the expected gains from trade—when markets are relatively competitive. As a result, initiatives to make these markets more competitive or transparent are only welfare-improving when both frictions are relatively severe, i.e., when buyers have a lot of market power (i.e., when π is low) and the adverse selection problem is relatively severe (i.e., when ϕ_1 is high).

While these comparative statics are certainly informative, one may be concerned that they reflect an inefficiency in the particular game we postulate between buyers and sellers. At the end of this section, we compare equilibrium outcomes to a constrained efficient benchmark. We argue that, in the region of the parameter space where $\phi_1 > 0$, equilibrium gains from trade coincide with those in the constrained efficient benchmark. This suggests that the comparative statics results that we described above are not a consequence of the particular game we have modeled, but rather a more fundamental feature of markets with adverse selection and imperfect competition.

²⁹Using price dispersion to help identify search frictions is standard in the industrial organization literature; see, e.g., ?.

5.1 Utilitarian Welfare

As noted above, our metric for economic activity will be the objective of a utilitarian planner, defined as the expected gains from trade realized between buyers and sellers, or

$$W(\pi, \mu_h) = (1 - \mu_h)(v_l - c_l) + \mu_h \left\{ \frac{2 - 2\pi}{2 - \pi} \int [x_h(u_l)(v_h - c_h)] dF_l(u_l) + \frac{\pi}{2 - \pi} \int \mu_h [x_h(u_l)(v_h - c_h)] d(F_l(u_l)^2) \right\}, \quad (19)$$

where, in a slight abuse of notation, we let

$$x_h(u_l) = 1 - \frac{u_h(u_l) - u_l}{c_h - c_l}. \quad (20)$$

The first term in (19) represents the gains from trade generated by low-quality goods; since all sellers receive at least one offer and $x_l = 1$ in every trade, all low-quality goods are transferred to the buyer. The second term captures the expected gains from trade between buyers and captive high-quality sellers. In particular, from equation (1), we can write the measure of captive sellers as $1 - p = \frac{2 - 2\pi}{2 - \pi}$. A randomly selected captive high-quality seller transfers $x_h(u_l)$ to the buyer and consumes the remaining $1 - x_h(u_l)$ herself, where u_l is drawn from $F_l(u_l)$. Finally, the last term in (19) captures the expected gains from trade between buyers and noncaptive high-quality sellers. A measure $p = \frac{\pi}{2 - \pi}$ of sellers are noncaptive and, since noncaptive sellers choose the maximum indirect utility among the two offers they receive, they trade an amount $x_h(u_l)$, where u_l is drawn from $F_l(u_l)^2$.

5.2 Increasing Competition

We first study the effects of increasing competition, which has been a common policy response to address perceived failures in markets for insurance, credit, and certain types of financial securities.³⁰ We do so by examining the relationship between welfare and competition, as captured by π . In Proposition 5, we establish that welfare is maximized at $\pi = 0$ when the adverse selection problem is relatively mild. However, when the adverse selection problem is severe, we show that W is hump-shaped in π ; i.e., there

³⁰For example, a recent report by the ? argues for “fostering greater competition” in health insurance plans by developing “policies that would increase the average number of sponsors per region,” which would then “increase the likelihood that beneficiaries would select low-cost plans.” Similarly, ? argued that the Consumer Financial Protection Bureau “will make consumer financial markets more transparent – and that’s good for everyone: The agency will give Americans [...] the tools they need to comparison shop for the best prices and the best loans, which will [...] increase competition and innovations that benefit borrowers.” A similar rationale underlies the Core Principles and Other Requirements for Swap Execution Facilities (?), issued under the Dodd-Frank Wall Street Reform and Consumer Protection Act, which requires that a swap facility sends a buyer’s request for price quotes to a minimum number of sellers before a trade can be executed.

is an *interior* level of competition that maximizes welfare in this region of the parameter space.

Proposition 5. *If $\phi_l \leq 0$, welfare is maximized at $\pi = 0$. Otherwise, it is maximized at a $\pi \in (0, 1)$.*

The first result in Proposition 5 is straightforward. Since a monopsonist offers a pooling contract in this region of the parameter space, all gains from trade are realized. Competition only serves to increase incentives to cream-skim. When these incentives are sufficiently strong, equilibrium menus offer high-quality sellers a higher price but a lower quantity to trade in order to ensure that such a deviation is not profitable, causing a decline in welfare.

The second result is less obvious. To see the intuition, first note that, as π increases, buyers allocate more surplus to sellers: $F_l(u_l)$ shifts to the right (in the sense of first-order stochastic dominance) and \bar{u}_l increases. Second, and crucially, $x_h(u_l)$ is hump-shaped in u_l : it increases near the monopsony offer c_l , and it decreases when u_l is sufficiently close to the competitive offer, v_l . When π is close to zero, \bar{u}_l is relatively small, and the distribution of offers is clustered near the monopsony contract; a small increase in π causes a rightward shift in the density of offers to values of u_l associated with *higher* values of x_h , increasing the gains from trade realized between buyers and high-quality sellers. When π is close to 1, \bar{u}_l is close to v_l , and the distribution of offers is clustered near the competitive contract; a small increase in π causes a shift toward values of u_l associated with *lower* values of x_h .

Therefore, understanding why welfare is hump-shaped in π ultimately requires understanding why $x_h(u_l)$ is hump-shaped in u_l . Note that, *ceteris paribus*, an increase in u_l relaxes the type l seller's incentive compatibility constraint, allowing buyers to raise x_h . In contrast, *ceteris paribus*, an increase in u_h tightens the type l seller's incentive compatibility constraint, requiring buyers to lower x_h . Thus, as offers to both types increase, the net effect on x_h depends on which one rises faster—formally, whether $U'_h(u_l)$ is greater or less than 1. Figure 4 illustrates this relationship between the quantity traded with high types, x_h , and the rate at which u_h and u_l increase within the set of equilibrium menus being offered. The figure reveals that u_l rises faster than u_h for smaller values of u_l , so that x_h is increasing in this region. However, as u_l nears v_l , u_h rises faster, and thus x_h is decreasing in this region.

To explain the hump-shape of welfare, then, we need to understand why $U'_h(u_l) < 1$ for low levels of u_l and $U'_h(u_l) > 1$ for high levels of u_l . While this slope is a complicated equilibrium object, determined by the interaction of an individual buyer's optimal strategy and the equilibrium distribution of offers, the basic intuition can be understood through two opposing forces. First, it is cheaper for buyers to provide utility to the low type (relative to the high type) because doing so has the additional benefit of relaxing the incentive constraints; we call this the “incentive effect,” and this force tends to reduce the

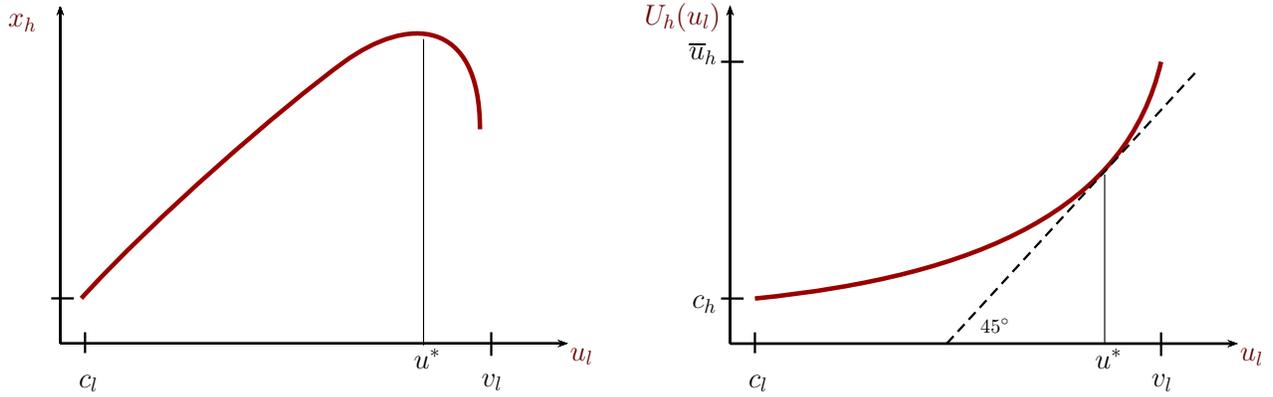


Figure 4: Trade (x_h) and Utility (U_h) of high-quality seller as functions of u_l when $\pi > 0$.

slope, $U'_h(u_l)$. Second, as u_l rises, buyers have more incentive to attract type h sellers, relative to type l sellers; formally, one can show that $\Pi_h(u_l, U_h(u_l))/\Pi_l(u_l, u_h)$ is increasing in u_l . This effect, which we call the “composition effect,” leads them to increase u_h at faster rates at higher u_l .

To illustrate these two forces more clearly, consider the following optimality condition that any equilibrium menu $(u_l, U_h(u_l))$ must satisfy:³¹

$$U'_h(u_l) = \underbrace{\frac{\phi_l}{\phi_h}}_{\text{incentive effect}} \underbrace{\frac{\Pi_h(u_l, U_h(u_l))}{\Pi_l(u_l)}}_{\text{composition effect}}, \quad (21)$$

where $\phi_h = (v_h - c_l)/(c_h - c_l)$ is the marginal cost of providing an additional unit of utility to type h sellers—i.e., $\phi_h = \frac{d\Pi_h}{du_h}$ —and for notational convenience $\Pi_l(u_l) \equiv \Pi_l(u_l, u_h)$. The first term, the incentive effect, is the ratio of the marginal costs of providing utility to the two types of sellers. Since this term is strictly less than 1, all else equal, the incentive effect leads to more aggressive competition for the low type and, therefore, to u_h rising more slowly than u_l .

The second term, the ratio of profits, can be larger or smaller than 1, depending on u_l . When u_l is close to the monopsony outcome, $\Pi_h \approx 0$, so the composition effect is also less than 1, and we have $U'_h(u_l) < 1$. However, as u_l approaches the upper bound v_l , this second term overwhelms the incentive effect, resulting in $U'_h(u_l) > 1$. In fact, one can show that $\lim_{u_l \rightarrow v_l} \frac{\Pi_h}{\Pi_l} = \lim_{u_l \rightarrow v_l} U'_h = \infty$.³² This is

³¹This equation combines the optimality condition (16) for u_l , the corresponding optimality condition for u_h , $\frac{\pi f_h}{1-\pi+\pi F_h} \Pi_h = \phi_h$, and the strict rank-preserving property $F_l(u_l) = F_h(U_h(u_l))$, which implies $f_l = f_h U'_h(u_l)$.

³²Indeed, one can show that, if this limit were finite, the equality in (21) cannot hold: the right-hand side must be strictly less than the left-hand side. Intuitively, if the ratio of profits is finite in the limit, buyers have incentive to offer a lower u_h . The only way to discourage such deviations is to make high types more profitable—in the limit, infinitely so. See the working paper version, ?, for a detailed discussion.

why $x'_h(u_l) < 0$ close to the Bertrand outcome.

5.3 Reducing Information Asymmetries

We now study the welfare consequences of reducing informational asymmetries. This exercise sheds light on the implications of certain policy initiatives, as well as the effects of various technological innovations. For example, an important debate in insurance, credit, and financial markets centers around information that the informed party (the seller in our context) is required to disclose and the extent to which such information can be used by the uninformed party (the buyers in our model) to discriminate.³³ Moreover, technological developments in these markets also have the potential to decrease informational asymmetries, as advanced record-keeping and more sophisticated scoring systems (e.g., credit scores) provide buyers with more and/or better information about sellers' intrinsic types.

To study such changes, we introduce a noisy *public* signal $s \in \{0, 1\}$ about the quality of each seller.³⁴ The signal is informative, so that $\Pr(s = 1|h) = \Pr(s = 0|l) > 0.5$. Since the signal is public, the buyers may condition their offers on it, i.e., offer separate menus for sellers with $s = 0$ and $s = 1$. Thus, the economy has two subgroups, $j \in \{0, 1\}$, with the fraction of high-quality sellers in subgroup j given by

$$\mu_{hj} = \frac{\mu_h \Pr(s = j | h)}{\mu_h \Pr(s = j | h) + \mu_l (1 - \Pr(s = j | l))}.$$

Note that the average across subgroups is equal to the unconditional fraction of high types, i.e., $\mathbb{E}[\mu_{hj}] = \mu_h$. The equilibrium outcome for each subgroup can be constructed using the procedure in Section 4 with the appropriate μ_{hj} . Welfare is then given by the average welfare across subgroups, i.e., $\mathbb{E}[W(\pi, \mu_{hj})]$. When buyers do not observe a signal (or, equivalently, are not permitted to condition their offers on it), welfare is simply $W(\pi, \mathbb{E}[\mu_{hj}])$. Hence, whether the signal increases or decreases welfare, respectively, depends on whether $W(\pi, \mu_h)$ is convex or concave in μ_h in the relevant region.

Before proceeding, two comments are in order. First, our focus is on the effect of a *small* increase in the information available to buyers; that is, we are interested in signals that induce a local mean-

³³In insurance markets, these questions typically concern an individual's health factors, both observable (e.g., age or gender) and unobservable (e.g., pre-existing conditions) to the insurance provider. In credit markets, similar questions arise with respect to observable characteristics that can legally be used in determining a borrower's creditworthiness, as well as the amount of information about a borrower's credit history that should be available to lenders (e.g., how long a delinquency stays on an individual's credit history). In financial markets, the relevant issue is not only whether a seller discloses relevant information about an asset to a buyer, but also whether the payoff structure of the asset is sufficiently *transparent* for sellers to distinguish good from bad assets. For example, in order to support "sustainable securitisation markets," the Basel Committee on Banking Supervision and the International Organization of Securities Commissions established a joint task force to identify criteria for "simple, transparent, and comparable" securitized assets. See <http://www.bis.org/bcbs/publ/d304.pdf>.

³⁴The restriction to a binary signal is only for simplicity. It is easy to introduce richer information structures.

preserving spread around μ_h . Very informative signals always improve welfare—for example, if buyers receive a perfect signal about sellers’ types, then all gains from trade are realized—but this is not a very interesting or realistic experiment. Second, we focus on the region with $\mu_h < \mu_0$, so that $\phi_l > 0$, which is more tractable and shows interesting interactions between competition and additional information.³⁵ Moreover, in this region, W is *linear* in μ_h when $\pi = 0$ or $\pi = 1$. Hence, imposing monopsony or perfect competition would lead us to the conclusion that additional information has no effect on welfare.

Proposition 6 shows that W has a strictly convex region when π is sufficiently low, implying that more information is beneficial when markets are close to (but not at) the monopsony benchmark. Alternatively, when markets are relatively (but not perfectly) competitive, W has a strictly concave region, implying that more information actually reduces welfare.

Proposition 6. *There exist $\underline{\pi}, \bar{\pi} \in (0, 1)$ such that: (i) for all $\pi \in (0, \underline{\pi})$, there exists $0 < \underline{\mu}_h < \bar{\mu}_h < \mu_0$ such that W is strictly convex on the interval $[\underline{\mu}_h, \bar{\mu}_h]$; and (ii) for all $\pi \in (\bar{\pi}, 1)$, there exists $0 < \underline{\mu}'_h < \bar{\mu}'_h < \mu_0$ such that W is strictly concave on the interval $[\underline{\mu}'_h, \bar{\mu}'_h]$.*

To see the intuition behind Proposition 6, recall from the previous subsection that trade with the high-quality seller (and thus welfare) is governed by the interaction of the incentive effect and the relative profit (or composition) effect. The consequences of more information can be understood in terms of these two forces, too. In particular, a lower ϕ_l drives down the first term in (21), which encourages more competition for low-quality sellers and, hence, boosts trade and welfare. Now, from (13), we see that ϕ_l is a concave function of μ_h . Since the additional signal induces a mean-preserving spread of μ_h , it results in a lower ϕ_l on average, which, *ceteris paribus*, increases trade. This mechanism makes more information desirable. The effect from relative profits goes in the opposite direction. In equilibrium, milder adverse selection raises profits from high types relative to low types, which increases U'_h and hence decreases trade. Close to monopsony, since the incentive effect dominates, more information raises welfare. The opposite happens when π is close to 1, and the effect on relative profits dominates.

5.4 Constrained Efficiency

The analysis above establishes that, when $\phi_l > 0$, increasing competition and reducing information asymmetries can have nonmonotonic effects on the equilibrium volume of trade, and hence on the (utilitarian) welfare measure. However, one might worry that this nonmonotonicity is an artifact of the particular game we study, and not a robust feature of markets with asymmetric information. To

³⁵Numerical simulations suggest that additional information always reduces welfare when $\mu_h > \mu_0$.

address this concern, in Appendix B, we adopt a mechanism design approach to derive a constrained efficient benchmark. A direct mechanism prescribes transfers of goods and numeraire for each buyer and seller based on their reported types, which include the quality of their good and the other agents with whom they are matched. An allocation is constrained efficient if it is implementable by a direct mechanism and maximizes a Pareto-weighted sum of utilities, subject to feasibility (only matched agents can trade), individual rationality (each agent receives a payoff at least as high as in equilibrium), incentive compatibility (types are reported truthfully), and exclusivity (sellers can trade with at most one buyer).

A key result is that the expected volume of trade in this benchmark coincides with that in our equilibrium when $\phi_l > 0$. This implies that a benevolent planner cannot propose a trading mechanism that would strictly increase trading volume or welfare, relative to our equilibrium allocation, given $\phi_l > 0$ and $\pi \in [0, 1]$. In contrast, when $\phi_l \in [\phi_2, 0]$ —where ϕ_2 is defined in Proposition 4—then the constrained efficient allocation implies full pooling ($x_h = 1$) in all trades while the expected trading volume of high-quality goods in equilibrium is strictly less than 1. In this region, a benevolent planner could make everyone better off by inducing more cross-subsidization from high- to low-quality sellers. The source of this inefficiency is similar to that which arises in many models with adverse selection and competition (see, e.g., ??): buyers’ incentives to cream-skim high-quality sellers limit equilibrium cross-subsidization.³⁶

Finally, another natural question is whether the welfare-maximizing allocation is attained in an environment where the level of competition is determined endogenously by the choices of market participants. In the next section, we take up this question by extending our analysis to study an environment where the market structure—summarized by π —is endogenous.³⁷

³⁶Hence, as in many models of adverse selection (i.e., hidden information about the common value of the asset), policies that incentivize pooling can potentially improve welfare when the equilibrium is constrained inefficient. In the context of our model, these policies would include restrictions on the type of contracts that buyers offer (e.g., ruling out separating contracts) or mandates that sellers must trade a minimum quantity. Note, however, that if sellers had hidden information about their private valuation, one may obtain constrained efficiency in an environment with directed search, implying no role for policy interventions (see, e.g., ?).

³⁷A related exercise is to consider interventions that mimic the effects of increasing or decreasing competition. In ?, we study what happens when the government enters a market suffering from adverse selection as a “large buyer,” as it has in, e.g., the markets for student loans, health insurance, or certain financial assets. We show that, by offering to buy any quantity at a fixed price, the government can increase sellers’ outside option and promote more competition, which recreates the effects of increasing π . Such an intervention can increase welfare only when both market power and the distortions arising from adverse selection are severe. Otherwise, we show that such programs can be detrimental to welfare even if, in principle, the intervention makes non-negative profits.

6 Endogenous Market Structure

In this section, we allow buyers to choose how intensely they advertise their offers to sellers. This exercise has two benefits: the *degree of competition* will be endogenous; and the measure of sellers who are contacted by at least one buyer, or *coverage*, will also be endogenous. This allows us to study which features of the environment determine the market structure and the corresponding welfare implications.

To this end, suppose that, in addition to choosing a menu of contracts to offer, each buyer $k \in \{1, 2\}$ also chooses the effort or intensity with which his offer will be advertised to sellers. Exerting effort is costly, but increases the likelihood that each seller observes the offer. Formally, we assume that buyer k can choose the probability $\hat{\pi}^k$ that each seller observes his offer by incurring a cost $C(\hat{\pi}^k)$, which is a continuously differentiable, strictly increasing, and strictly convex function with $C(0) = C'(0) = 0$ and $C'(1) = \infty$.³⁸ Note that $\hat{\pi}^k$ represents a slightly different object than π in our benchmark model, since it affects both competition and coverage. However, what is crucial is that—just like π in our earlier analysis— $\hat{\pi}^{-k}$ is the conditional probability that a seller who buyer k meets has a second offer. Hence, in a symmetric Nash equilibrium, $\hat{\pi}^k = \hat{\pi}^{-k} \equiv \hat{\pi}$ remains the key determinant of the level of competition.

Taking as given the other buyer's advertising intensity, $\hat{\pi}^{-k}$, and the distribution of offers that he makes to sellers of type $i \in \{l, h\}$, $F_i^{-k}(u_i^{-k})$, buyer k chooses a tuple $(\hat{\pi}^k, u_l^k, u_h^k)$ to maximize

$$\sum_{i \in \{l, h\}} \mu_i [\hat{\pi}^k (1 - \hat{\pi}^{-k}) + \hat{\pi}^k \hat{\pi}^{-k} F_i^{-k}(u_i^k)] \Pi_i(u_l^k, u_h^k), \quad (22)$$

subject to the participation and incentive constraints described in the benchmark model, with $\Pi_i(\cdot, \cdot)$ defined in (9)–(10). Factoring out $\hat{\pi}^k$ from (22), one can see that the choice of $\hat{\pi}^k$ and (u_l^k, u_h^k) are separable. Hence, given $\hat{\pi}^{-k}$, the first-order conditions on u_l^k and u_h^k are exactly as they were before (replacing π with $\hat{\pi}^{-k}$), while the first-order condition determining the optimal choice of $\hat{\pi}^k$ is

$$C'(\hat{\pi}^k) = \sum_{i \in \{l, h\}} \mu_i [1 - \hat{\pi}^{-k} + \hat{\pi}^{-k} F_i^{-k}(u_i^k)] \Pi_i(u_l^k, u_h^k). \quad (23)$$

In a symmetric equilibrium, where $\hat{\pi}^1 = \hat{\pi}^2 \equiv \hat{\pi}$, equation (23) implies that the marginal cost of increasing $\hat{\pi}$ is equal to the equilibrium profits characterized in Propositions 2 and 3. Since these profits are decreasing in $\hat{\pi}$, the next result follows almost immediately.

³⁸Note that this implies a fraction $(1 - \hat{\pi}^1)(1 - \hat{\pi}^2)$ of sellers receive zero offers. This is the sense in which coverage is endogenous in the current setup, whereas we fixed this fraction (to zero) in our benchmark model.

Proposition 7. For any $\phi_l < 1$, there exists a unique symmetric equilibrium, with $\hat{\pi}^* \in (0, 1)$ and $\{F_i^*(u_i)\}_{i \in \{l, h\}}$ as described in Propositions 2 and 3.

In Lemma 5, below, we offer comparative statics with respect to the fraction of high-quality sellers, μ_h .³⁹ Recall that there exists a μ_0 such that $\phi_l \geq 0$ if and only if $\mu_h \leq \mu_0$. We show that the equilibrium $\hat{\pi}^*$ is U-shaped in μ_h , achieving a minimum at $\mu_h = \mu_0$.

Lemma 5. The equilibrium advertising intensity $\hat{\pi}^*$ is decreasing in μ_h if $\mu_h < \mu_0$ and increasing otherwise.

To understand the intuition, consider first the case of “severe adverse selection,” i.e., when $\mu_h < \mu_0$ or, equivalently, when $\phi_l > 0$. In this region, once information rents are taken into account, the buyer’s payoff from trading with low-quality sellers is larger than the payoff from trading with high-quality sellers (even if $v_h - c_h > v_l - c_l$). Thus, from the buyer’s perspective, an increase in μ_h in this region actually worsens the pool of potential sellers and, as a result, buyers optimally choose a lower $\hat{\pi}$. The opposite is true when $\mu_h > \mu_0$, where we say adverse selection is “mild.” In this region, after adjusting for information rents, it is relatively more profitable to trade with high-quality sellers, and thus buyers optimally choose larger values of $\hat{\pi}$ as the fraction of high-quality sellers increases.

Lemma 5 has implications for the relationship between the composition of high- and low-quality sellers in a market and the (endogenous) level of competition that prevails. In particular, this result suggests that competition for customers should be strongest in markets with less uncertainty about sellers’ types (i.e., extreme values of μ_h), and weakest in markets with more uncertainty about sellers’ types (i.e., intermediate values of μ_h).⁴⁰

The model with endogenous $\hat{\pi}$ also allows us to connect some of our welfare results to more concrete implications for policy. To see this, suppose $C(\hat{\pi}^k) = A c(\hat{\pi}^k)$ for some positive constant $A > 0$, and consider the effect of taxing buyers’ advertising intensities according to a proportional tax, $\tau \hat{\pi}$. For simplicity, suppose all tax proceeds are then simply rebated to the agents. In Lemma 6, below, we establish that welfare is increasing in τ in some regions of the parameter space. That is, a policy making it *more costly* for buyers to contact sellers can *improve* welfare.

Lemma 6. Suppose $\phi_l > 0$. There exists an $\tilde{A} > 0$ such that welfare is increasing in τ for all $A < \tilde{A}$.

The result in Lemma 6 follows closely from the fact that welfare is hump-shaped in $\hat{\pi}$, even taking

³⁹Comparative statics for other parameters may be similarly derived by considering the effect of perturbing each parameter on equilibrium profits.

⁴⁰One can also show that there exist cost functions that are consistent with *any* $\hat{\pi}^*$, so that endogenizing competition does not rule out certain types of equilibrium.

into account that an increase in $\hat{\tau}$ increases coverage. As a result, when A is sufficiently small, $\hat{\tau}^*$ is large and a decrease in $\hat{\tau}^*$ —brought about by an increase in τ —causes welfare to rise.

7 Large Markets and Meeting Technologies

We now show how our analysis and results extend to an environment with an arbitrarily large number of buyers and sellers and a more general *meeting technology*. Suppose there is a measure b of buyers and a measure s of sellers. As in our benchmark model, buyers send offers and sellers receive them. The meeting technology dictates the number of offers each buyer sends and where these offers end up.

Formally, let η denote the (expected) number of offers that each buyer sends; let $\lambda = \frac{\eta b}{s}$ denote the ratio of offers to sellers; and let P_n denote the probability that each seller receives $n \in \{0\} \cup \mathbb{N}$ offers. A meeting technology, then, can be summarized by a pair (λ, P_n) .⁴¹ For a buyer, a meeting technology implies that an offer he sends is received by a seller with $n - 1$ other offers with probability Q_n , where $nP_n = \lambda Q_n$ for all $n \in \mathbb{N}$. Following the convention in the literature, we let $Q_0 = 1 - \sum_{n=1}^{\infty} Q_n$.

7.1 Characterizing Equilibrium

As in our benchmark model, we restrict attention to symmetric equilibria, where $\{F_i(u_i)\}_{i \in \{l, h\}}$ summarizes the distribution of menus being offered by buyers. Taking this distribution as given, an individual buyer makes an offer (u_l, u_h) that solves

$$\max_{u_l, u_h} \sum_{i \in \{l, h\}} \mu_i \left[\sum_{n=1}^{\infty} Q_n F_i^{n-1}(u_i) \right] \Pi_i(u_l, u_h), \quad (24)$$

where, again, $\Pi_i(u_l, u_h)$ is defined in (9)–(10). Importantly, the objective in (24) can be rewritten as

$$[1 - Q_0] \sum_{i \in \{l, h\}} \mu_i [1 - \tilde{\pi} + \tilde{\pi} G_i(u_i)] \Pi_i(u_l, u_h), \quad (25)$$

where $\tilde{\pi} = 1 - Q_1/(1 - Q_0)$ is the probability that an offer is received by a seller that has at least one other offer, conditional on being received by a seller, and

$$G_i(u_i) = \frac{1}{\tilde{\pi}} \sum_{n=2}^{\infty} \frac{Q_n}{1 - Q_0} F_i^{n-1}(u_i) \quad (26)$$

⁴¹This formulation of a meeting technology is slightly more general than what is commonly used in the existing literature (e.g., ?), in that we allow the “queue length” λ to depend on the meeting technology.

is the probability that the seller accepts the offer u_i , given that they own a good of quality $i \in \{l, h\}$.

Notice immediately that (25) has the same form as our objective function in the two-buyer case, replacing π with $\tilde{\pi}$ and $F_i(u_i)$ with $G_i(u_i)$. As a result, our characterization of equilibrium in Propositions 2 and 3 is preserved, and the distribution $G_i(u_i)$ is uniquely defined in all regions of the parameter space. Moreover, from (26), it is easy to show that $G_i(u_i)$ uniquely determines the distribution of offers made by buyers, $F_i(u_i)$. Using these results, one can easily determine the type of contracts that are offered in equilibrium⁴²; the distribution of offers that are made to each type of seller, $F_i(u_i)$, which is the solution to (26); and the prices and quantities that are ultimately traded in equilibrium.

7.2 Competition, Coverage, and Equilibrium Gains from Trade

In our benchmark model, we studied the effects of changing the probability that a seller received two offers, π . We now explore similar comparative statics within the context of a general meeting technology. In particular, we let P_n and λ (and hence Q_n) depend on a parameter α . This formulation is intentionally general: a change in α could correspond to a change in the measure of buyers, a change in the expected number of offers per buyer, or a change in the technology that matches offers to sellers.

As in Section 6, we focus on the case where $\phi_l > 0$ and define the utilitarian welfare measure

$$W(\alpha) = \sum_{n=1}^{\infty} P_n(\alpha) \left[\mu_h (v_h - c_h) \int x_h(u_l) d(F_l^n(u_l)) + \mu_l (v_l - c_l) \right] + \sum_{i=l,h} \mu_i c_i.$$

As in our benchmark model, when $\phi_l > 0$, the distribution $G_l(u_l)$ solves the differential equation

$$\frac{\tilde{\pi} g_l(u_l)}{1 - \tilde{\pi} + \tilde{\pi} G_l(u_l)} = \frac{\phi_l}{v_l - u_l} \quad (27)$$

with support $[c_l, \bar{u}_l(\alpha)]$ such that $G_l(c_l) = 0$ and $G_l(\bar{u}_l(\alpha)) = 1$.

Solving (27) and imposing equal profits implies that the mapping $U_h(u_l)$ must satisfy

$$\left(\frac{v_l - c_l}{v_l - u_l} \right)^{\phi_l} [\mu_l (v_l - u_l) + \mu_h \Pi_h(u_l, U_h(u_l))] = \mu_l (v_l - c_l),$$

exactly as in the case of two buyers. An immediate, and important, implication is that $x_h(u_l)$ is hump-shaped in u_l and independent of α . Hence, a change in α only affects the distribution of offers that are made, summarized by F_l , and the distribution of offers that sellers receive, summarized by P_n .

⁴²This is done by comparing ϕ_l , which is unchanged, to ϕ_1 and ϕ_2 , which are updated by replacing π with $\tilde{\pi}$.

As a result, the effects of a change in α can be decomposed as follows:

$$W'(\alpha) = \sum_{n=1}^{\infty} \frac{\partial P_n(\alpha)}{\partial \alpha} \left[\mu_h (v_h - c_h) \int x_h(u_l) d(F_l^n(u_l; \alpha)) + \mu_l (v_l - c_l) \right] + \sum_{n=1}^{\infty} P_n(\alpha) \left[\mu_h (v_h - c_h) \frac{\partial}{\partial \alpha} \int_{c_l}^{\bar{u}_l(\alpha)} x_h(u_l) d(F_l^n(u_l; \alpha)) \right],$$

where, for the purpose of clarity, we've made the dependence of F_l on α explicit. The first term in the equation above was absent in our benchmark model, but captures a standard effect in models with frictions: the effect of a change in α on the set of sellers who are able to trade, or what we call the *coverage* effect. The second term captures the effect that we focused on in our benchmark model: the effect of a change in α on the distribution of offers, or what we call the *competition* effect.

For example, suppose increasing α leads to a first-order stochastic dominant (FOSD) shift in the number of offers that sellers receive. In this case, the coverage effect would be positive, since fewer sellers receive zero offers. However, the competition effect could be negative, since an increase in α leads to a FOSD shift in the distribution of offers F_l . As in our benchmark model, when α is sufficiently large, this shift puts more weight on the downward-sloping region of $x_h(u_l)$, thus reducing welfare.

Which of these two effects dominates typically depends on the details of the meeting technology. Consider, for example, the Poisson meeting technology with $\lambda(\alpha) = \alpha$ and $P_n(\alpha) = \frac{e^{-\alpha} \alpha^n}{n!}$. This is perhaps the most popular meeting technology in the literature (see, e.g., ? and ? for early examples, and ? for a more recent example), and an increase in α clearly leads to a FOSD shift in the distribution of offers that sellers receive. We show in the Appendix that when $\phi_l > 0$, there exists an α^* such that welfare is decreasing in α for all finite $\alpha > \alpha^*$. Therefore, as in our benchmark model, some frictions can increase welfare *even after* accounting for the coverage effect.

The same is not true, however, for all meeting technologies. For example, consider the Geometric meeting technology with $\lambda(\alpha) = \alpha/(1-\alpha)$ and $P_n(\alpha) = \alpha^n(1-\alpha)$, which was studied recently by, e.g., ?. Under this meeting technology, when $\phi_l > 0$, we show in the Appendix that the coverage effect always dominates the competition effect, so that welfare is increasing in α . Intuitively, the coverage effect is relatively strong because the fraction of sellers who fail to receive an offer, P_0 , falls slowly in α as $\alpha \rightarrow 1$, whereas the competition effect vanishes more quickly.⁴³

⁴³Note, however, that one can augment the Geometric meeting technology to ensure full coverage by setting $\lambda(\alpha) = \alpha/(1-\alpha)$ and $P_n(\alpha) = \alpha^{n-1}(1-\alpha)$ for $n \in \mathbb{N}$, with $P_0 = 0$. This specification removes the positive effects of increased coverage on welfare, leaving only the negative competition effect as $\alpha \rightarrow 1$.

8 Conclusion

In their survey of the literature on insurance markets, ? note that, despite substantial progress in understanding the effects of adverse selection,

“There has been much less progress on empirical models of insurance market competition, or on empirical models of insurance contracting that incorporate realistic market frictions. One challenge is to develop an appropriate conceptual framework. Even in stylized models of insurance markets with asymmetric information, characterizing competitive equilibrium can be challenging, and the challenge is compounded if one wants to allow for realistic consumer heterogeneity and market imperfections.”

In this paper, we overcome this challenge and develop a tractable, unified framework to study adverse selection, screening, and imperfect competition. We provide a full analytical characterization of the unique equilibrium, and use it to study both positive and normative issues.

Our framework can be exploited and extended to address a variety of important issues. On the applied side, our equilibrium provides a new structural framework that can be used to jointly identify the extent of adverse selection and imperfect competition in various markets, and to study how the interaction of these two frictions affects the distribution of contracts, prices, and quantities that are traded. On the theoretical side, one natural extension is to study the analog of our model with *nonexclusive contracts*; although this would complicate the analysis considerably, it would also make our framework suitable to analyze certain markets where exclusivity is hard to enforce. We leave these exercises for future work.

References