EFRSW
7/15/08
4:30 pm

Competitive Search with Private Information

Veronica Guerrieri
University of Chicago

Robert Shimer
University of Chicago

Randall Wright
University of Pennsylvania

July 9, 2008
Preliminary

1 Introduction

This paper studies the interaction between a large number of uninformed principals and a large number of informed agents in a market. Principals compete for agents by posting contracts at some cost. Agents observe the posted contracts and try to obtain one of them. Their ability to obtain a contract may be limited by search frictions or capacity constraints. For example, if each principal can match with at most one agent and fewer principals offer a particular contract than the number of agents who wish to obtain it, each agent is served only probabilistically. If an agent obtains a contract, the contract fully specifies the terms of trade and hence the principal's and agent's expected payoff.

Part of the contribution of this paper is technical: we develop a canonical extension to the competitive search model (Montgomery, 1991; Peters, 1991; Moen, 1997; Shimer, 1996) that allows for private information.[1] We prove that in any equilibrium, principals offer separating contracts, so different types of agents try to obtain different contracts. We use this to prove existence of a competitive search equilibrium with private information and show that, if multiple equilibria exist, they are Pareto ranked. On the other hand, separation is socially costly, especially when there are relatively few of the undesirable types in the economy. A pooling allocation may Pareto dominate the equilibrium.

We further develop our model through a series of extended examples. The first shows that an equilibrium may not exist if the principals do not have enough instruments to screen the desirable types from the undesirable ones. In the paper, we make assumptions that ensure that there is a technology to separate the agents.

[1] See also Faig and Jerez (2005), Guerrieri (2008), and Moen and Rosén (2006).

The second example shows how search frictions or capacity constraints resolve the nonexistence issue in Rothschild and Stiglitz (1976). In the original paper, the authors prove that any equilibrium must have principals offering the least cost separating contract, with undesirable agents obtaining an unconstrained contract and desirable agents obtaining a distorted contract that does not attract the undesirables. Rothschild and Stiglitz (1976) find that, when there are relatively few undesirable agents, an equilibrium may fail to exist. When there are relatively few undesirable agents, a principal can make more profits by offering an undistorted contract that attracts both types of agents, losing money on the undesirables but more than compensating with the desirables. In a competitive search model, this deviation is infeasible, regardless of the composition of the pool of agents, because of the search friction or capacity constraint. As more agents are attracted to this undistorted contract, each is less likely to be served. Eventually this deters some of the agents from trying to obtain the contract, and the more desirable agents are always the first to disappear because their alternative possibility—trying to obtain a distorted contract—is more attractive. This means that a deviating principal only attracts undesirable agents and loses money on each of them. Although an equilibrium always exists, a Pareto improvement is feasible if there are sufficiently few undesirable agents.

In this example, contracts are distorted ex post, after the desirable agent has matched with a principal. We develop a third example where the distortion occurs through the probability of exchange. Some agents are more eager to trade than others, but the principal prefers to trade with the agent who is less eager to trade. For example, some agents may hold counterfeit money that they are eager to unload, while principals gain more from trading with an agent holding genuine money. In this case, the equilibrium may involve only probabilistic meetings for the agent holding the desirable good, even if it is feasible for each principal and each agent to meet someone. The novelty here is that the probability of a meeting, rather than the probability of trade within a meeting, screens out the undesirable agents. In this example, we show that whenever an equilibrium exists, it is Pareto dominated by a pooling allocation.

2 Model

There is a measure 1 of risk-neutral, ex-ante heterogeneous agents. A fraction $\pi_s > 0$ of agents are of type $s \in \{1, 2, \ldots, S\}$. The type is the agent's private information. Moreover, there is a large measure of risk neutral ex-ante homogeneous principals.

We study a static model, with only a single opportunity to match. At the start of time, each principal can post a contract $\mathcal{C}$ by paying a cost k. The next paragraph describes contracts in more detail. Each agent observes the set of posted contracts and tries to obtain one of them. Then matching takes place. If the ratio of principals offering contract $\mathcal{C}$ to agents seeking that contract, hereafter the *market tightness*, is θ, each agent finds a principal with probability $\mu(\theta) \in [0,1]$ and each principal finds an agent with probability $\mu(\theta)/\theta \in [0,1]$, where μ is an increasing, concave function and $\mu(\theta)/\theta$ is decreasing in θ. Since all agents who apply for a contract are equally likely to succeed, the probability that when a principal finds an agent, her type is s, is equal to the fraction of agents seeking that contract whose type is s, say $p_s \in [0,1]$. Matched principals and agents implement the terms of the contract and obtain the implied payoffs. Principals and agents who do not match get their outside option, which we normalize to 0.

A contract $\mathcal{C}$ consists of two parts: a revelation mechanism $Y = \{y_1, y_2, \ldots, y_S\}$ that specifies the terms of trade between a principal and an agent who reports her type is $s \in \{1, 2, \ldots, S\}$; and a recommendation $\theta \geq 0$ and $P = \{p_1, p_2, \ldots, p_S\}$, where $p_s \geq 0$ and $\sum_s p_s = 1$. Here θ is the recommended market tightness associated with the contract and p_s is the recommended share of type s agents among the pool of agents seeking that contract. We assume $y_s \in \mathbb{Y}$ for all s, a compact set. We invoke the revelation principal and focus on contracts where agents truthfully reveal their type when they match with a principal and where recommendations are followed; of course it is necessary to restrict attention to incentive compatible contracts, which we define below. The expected profit for a principal offering a direct revelation contract $\mathcal{C} = \{Y, P, \theta\}$ is

$$\frac{\mu(\theta)}{\theta} \sum_{s=1}^{S} p_s v_s(y_s) - k, \tag{1}$$

where $v_s(y_s)$ is the principal's payoff from terms of trade y_s with agent s, continuous and bounded for all $y \in \mathbb{Y}$. The expected utility of a type-s agent who applies for this contract and reports her type is s' is

$$\mu(\theta) u_s(y_{s'}), \tag{2}$$

where $u_s(y_{s'})$ is a type-s agent's payoff from terms of trade $y_{s'}$, continuous and bounded for all $y \in \mathbb{Y}$. Note that the terms of trade y_s is itself typically a vector, for example specifying a transfer and an action by the agent.

Definition 1 *Let $\mathbb{C}$ denote the set of incentive-compatible contracts. A contract $\mathcal{C} = \{Y, P, \theta\}$ is incentive compatible if and only if $u_s(y_s) \geq u_s(y_{s'})$, for all s, s'.*

We are now in a position to define an equilibrium. In a competitive search equilibrium, all principals assume that they cannot affect the expected utility of a type-s agent, say $\bar{U}_s$.

If a particular incentive compatible contract $\mathcal{C} = \{Y, P, \theta\}$ offers a type-s agent less utility than $\bar{U}_s$, she will be unwilling to look for that contract. If it offers her more utility, type-s agents will apply in greater numbers than implied by the tightness θ, and so this cannot be an equilibrium. In order to attract type-s agents, a contract must offer exactly $\bar{U}_s$. More precisely:

Definition 2 *A Competitive Search Equilibrium is a vector $\bar{U} \in \mathbb{R}_+^S$ and a measure λ on $\mathbb{C}$ with support $\overline{\mathbb{C}}$ satisfying*

1. *for any contract $\mathcal{C} = \{Y, P, \theta\} \in \overline{\mathbb{C}}$,*
   (a) *principals'* ***free entry*** *requires $\dfrac{\mu(\theta)}{\theta} \sum_{s=1}^{S} p_s v_s(y_s) = k$ and*
   (b) *type-s agents'* ***optimal search*** *requires $\bar{U}_s \geq \mu(\theta) u_s(y_s)$ with equality if $p_s > 0$;*
2. *for any incentive compatible contract $\mathcal{C}' = \{Y', P', \theta'\} \in \mathbb{C}$, the following two conditions cannot both be satisfied:*
   (a) *principals earn positive profits: $\dfrac{\mu(\theta')}{\theta'} \sum_{s=1}^{S} p'_s v_s(y'_s) > k$, and*
   (b) *agents' search is optimal: for all s, $\bar{U}_s \geq \mu(\theta) u_s(y'_s)$ with equality if $p_s > 0$;*
3. *markets clear: for all s, $\dfrac{1}{\theta} \displaystyle\int p_s d\lambda(\{Y, P, \theta\}) \leq \pi_s$, with equality if $\bar{U}_s > 0$.*

Condition (1a) ensures that principals earn zero profits in equilibrium. Condition (1b) implies that agents are attracted to a contract only if it delivers utility $\bar{U}_s$ and that no contract delivers more utility. Condition (2) ensures that there is no other incentive compatible contract that a principal could offer and make positive profits. Here (2b) ensures that when a principal considers offering such a contract, it recognizes that agents will allocate their search so as to ensure that this contract is no more attractive than any other contract and will not search for it unless it is equally attractive. Finally, condition (3) ensures that type-s agents search for some contract, unless they are indifferent about matching in equilibrium, $\bar{U}_s = 0$.

For the most part, this definition of competitive search equilibrium is standard; see for example Acemoglu and Shimer (1999). The most novel piece is that we treat the recommendation $\{P, \theta\}$ as part of the contract $\mathcal{C}$. The literature previously would have treated a contract as a revelation mechanism, $\mathcal{C} = Y$, and substituted the recommendations with functions of the mechanism Y, say schedules $\Theta(Y)$ and $\Pi(Y)$. Principals and agents both anticipate that if a principal offers a contract Y, the associated market tightness will be

$\theta = \Theta(Y)$ and the associated share of each type of agent will be $P = \Pi(Y)$. This specification does not permit two principals to offer identical revelation mechanisms, with one intending to attract type-1 agents while the other intends to attract type-2 agents, possibly with a different market tightness. Making the recommendation part of the contract allows for this possibility. With only one type of agent, i.e. in most of the previous literature, this distinction is immaterial. In any case, the recommendation is only cheap talk. A principal cannot force type-s agents to search for his contract if they do not want to, nor can he keep them away if the contract would offer them more utility than $\bar{U}_s$.

3 Characterization

The first step in characterizing an equilibrium is showing that it is equivalent to the solution to a set of constrained optimization problems:

Proposition 1 *If a vector $\bar{U} \in \mathbb{R}_+^S$ and a measure λ on $\mathbb{C}$ with support $\overline{\mathbb{C}}$ is a Competitive Search Equilibrium, then any $\{Y, P, \theta\} \in \overline{\mathbb{C}}$ solves*

$$k = \max_{Y,P,\theta} \frac{\mu(\theta)}{\theta} \sum_{s=1}^{S} p_s v_s(y_s) \tag{P1}$$

$$\text{s.t. } u_s(y_s) \geq u_s(y_{s'}) \text{ for all } s, s',$$

$$\text{and } \bar{U}_s \geq \mu(\theta) u_s(y_s) \text{ with equality if } p_s > 0 \text{ for all } s.$$

Conversely, if there exists a vector $\bar{U}$ and a measure λ such that each $\{Y, P, \theta\} \in \overline{\mathbb{C}}$ solves problem (P1) *and market clearing holds, then there exists a Competitive Search Equilibrium $\{\bar{U}, \lambda\}$.*

Proof. First suppose $\{\bar{U}, \lambda\}$ is a competitive search equilibrium. Take any $\{Y, P, \theta\} \in \overline{\mathbb{C}}$. Part (1a) of the definition of competitive search equilibrium implies that the value of the objective function in problem (P1) must be k. $\{Y, P, \theta\} \in \overline{\mathbb{C}}$ ensures it satisfies the first constraint in problem (P1), while part (1b) ensures it satisfies the second constraint. Now take any other vector $\{Y', P', \theta'\}$ that satisfies the constraints in (P1). Since $u_s(y'_s) \geq u_s(y'_{s'})$ for all s and s', the alternative contract is incentive compatible, $\{Y', P', \theta'\} \in \mathbb{C}$. Moreover, the second constraint in problem (P1) implies that the alternative contract satisfies condition (2b). Therefore part (2a) of the definition of equilibrium implies $k \geq \frac{\mu(\theta')}{\theta'} \sum_{s=1}^{S} p'_s v_s(y'_s)$. Thus $\{Y, P, \theta\}$ delivers a higher value in problem (P1) than any other vector satisfying the constraints, i.e. it solves the constrained optimization problem (P1).

Conversely, suppose there exists a vector $\bar{U}$ and a measure λ such that each $\{Y, P, \theta\} \in \overline{\mathbb{C}}$ solves problem (P1) and market clearing holds. Part (1a) of the definition of equilibrium follows from the objective function in (P1), while (1b) follows from the second constraint. Moreover, if part (2) of the definition of equilibrium failed, there would be an alternative contract that satisfied the constraints in (P1) and delivered a higher value than k, a contradiction. ■

To further aid in finding a competitive search equilibrium, we next prove that without loss of generality we can restrict attention to equilibria where each contract is designed to attract only one type of agent. Moreover, in such an equilibrium, there is no need to ask an agent her type once she matches with the principal; each principal gives all agents the same terms of trade.

Lemma 1 *Suppose there exists an equilibrium $\{\bar{U}, \lambda\}$ with some contract $\mathcal{C} = \{Y, P, \theta\}$ with $d\lambda(\mathcal{C}) > 0$ and either $p_s, p_{s'} > 0$ or $y_s \neq y_{s'}$ for some $s \neq s'$. Let $\{\{y_s\}, 1_s, \theta\}$ denote a contract that offers all agents terms of trade y_s and attracts only type s agents ($p_s = 1$). Then there exists an equilibrium $\{\bar{U}, \tilde{\lambda}\}$ such that $d\tilde{\lambda}(\mathcal{C}) = 0$, $d\tilde{\lambda}(\{y_s\}, 1_s, \theta) = p_s d\lambda(\mathcal{C})$ for all s and s', and $d\tilde{\lambda}(\mathcal{C}') = d\lambda(\mathcal{C})$ otherwise.*

Proof. Consider an equilibrium $\{\bar{U}, \lambda\}$ with some contract $\mathcal{C} = \{Y, P, \theta\}$ such that $d\lambda(\mathcal{C}) > 0$ and $p_s, p_{s'} > 0$ for some $s \neq s'$. For this contract to solve problem (P1), it must be the case that $v_s(y_s) = v_{s'}(y_{s'})$. For example, if $v_s(y_s) > v_{s'}(y_{s'})$, the contract $\{Y, 1_s, \theta\}$ would yield a higher value of the objective function in (P1) without modifying any of the constraints.

Now consider the contract $\mathcal{C}_s \equiv \{\{y_s\}, 1_s, \theta\}$. The value of the objective function in (P1) is unchanged at $k = \frac{\mu(\theta)}{\theta} v_s(y_s)$. Moreover, the constraints are satisfied. The first constraint holds because $y_s = y_{s'}$ for all s'. The second constraint holds for s because $\bar{U}_s = \mu(\theta) u_s(y_s)$ from the fact that the original contract $\mathcal{C}$ solved (P1). It holds for arbitrary s' because $\bar{U}_{s'} \geq \mu(\theta) u_{s'}(y_{s'}) \geq \mu(\theta) u_{s'}(y_s)$; the first inequality holds because the original contract $\mathcal{C}$ satisfied the last constraint in (P1), while the second inequality holds because $\mathcal{C}$ satisfied the first constraint. Since $\mathcal{C}_s$ is feasible and payoff equivalent to $\mathcal{C}$, it solves problem (P1).

Finally, we construct a separating equilibrium $\{\bar{U}, \tilde{\lambda}\}$ where $d\tilde{\lambda}(\mathcal{C}_s) = p_s d\lambda(\mathcal{C})$ and $d\tilde{\lambda}(\mathcal{C}') = d\lambda(\mathcal{C}')$ for any other $\mathcal{C}'$. This $\{\bar{U}, \tilde{\lambda}\}$ is an equilibrium, given that the contracts posted solve problem (P1) and market clearing is satisfied. ■

This lemma allows us to restrict attention to equilibria where each contract is designed to attract only one type of agent. This means that a contract designed to attract type s is characterized by $p_s = 1$ and $p_{s'} = 0$ for all $s' \neq s$. Hence, with a slight abuse of notation, we can represent a contract designed for type s simply by a pair $\{y_s, \theta_s\}$, where y_s denotes the terms of trade for agents who obtain this contract and θ_s is the recommended market tightness. Note that the subscript s denotes the type of agents that are supposed to search for this contract.

According to problem (P1), for each s, the contract $\{y_s, \theta_s\}$ solves the following constrained maximization problem:

$$
\begin{aligned}
k = &\max_{y,\theta} \frac{\mu(\theta)}{\theta} v_s(y) \\
\text{s.t. } &\bar{U}_s = \mu(\theta) u_s(y) \\
\text{and } &\bar{U}_{s'} \geq \mu(\theta) u_{s'}(y) \text{ for all } s'.
\end{aligned} \tag{P2}
$$

Finding a competitive search equilibrium now amounts to finding a vector $\{\bar{U}_s\}$ consistent with problem (P2) for all s. It is easier to work with the dual of this problem:

Lemma 2 *Let the contract* $\{y_s, \theta_s\}$ *solve*

$$
\begin{aligned}
\bar{U}_s = &\max_{y,\theta} \mu(\theta) u_s(y) \\
\textit{s.t. } &\frac{\mu(\theta)}{\theta} v_s(y) = k \\
\textit{and } &\bar{U}_{s'} \geq \mu(\theta) u_{s'}(y) \textit{ for all } s' \neq s.
\end{aligned} \tag{P3}
$$

If the first constraint is binding, so $\bar{U}_s < \max_{y,\theta} \mu(\theta) u_s(y)$ *subject to* $\bar{U}_{s'} \geq \mu(\theta) u_{s'}(y)$*, this contract also solves problem* (P2).

Proof. Suppose $\{y_s, \theta_s\}$ solves problem (P2). Construct the Lagrangian

$$
L(y, \theta, \lambda) = \lambda_0 \frac{\mu(\theta)}{\theta} v_s(y) + \sum_{i=1}^{S} \lambda_i \big(\bar{U}_i - \mu(\theta) u_i(y)\big).
$$

The Fritz John version of Lagrange's theorem (see, for example, Simon and Blume, 1994) states that there are numbers $\lambda_0^*, \lambda_1^*, \ldots, \lambda_S^*$, not all zero, with $\lambda_i^* \geq 0$ for $i \neq s$ and $\lambda_0^* \in \{0, 1\}$, such that

$$
\frac{\partial L}{\partial y} = \frac{\partial L}{\partial \theta} = \frac{\partial L}{\partial \lambda_s} = 0, \quad \frac{\partial L}{\partial \lambda_i} \geq 0,\ i \neq 0, s, \quad \text{and } \lambda_i \frac{\partial L}{\partial \lambda_i} = 0,\ i \neq 0, s,
$$

when evaluated at $\{y_s, \theta_s, \lambda^*\}$.

Similarly, if $\{y_s, \theta_s\}$ solves problem (P3), construct the Lagrangian

$$
\tilde{L}(y, \theta, \tilde{\lambda}) = \tilde{\lambda}_s \mu(\theta) u_s(y) + \tilde{\lambda}_0 \left(\frac{\mu(\theta)}{\theta} v_s(y) - k \right) + \sum_{i \neq s, 0} \tilde{\lambda}_i \big(\bar{U}_i - \mu(\theta) u_i(y)\big).
$$

There are numbers $\tilde{\lambda}_0^*, \tilde{\lambda}_1^*, \ldots, \tilde{\lambda}_S^*$, not all zero, with $\tilde{\lambda}_i^* \geq 0$ for all $i \neq 0$ and $\tilde{\lambda}_s^* \in \{0, 1\}$, such that

$$\frac{\partial \tilde{L}}{\partial y} = \frac{\partial \tilde{L}}{\partial \theta} = \frac{\partial \tilde{L}}{\partial \tilde{\lambda}_0} = 0, \quad \frac{\partial \tilde{L}}{\partial \tilde{\lambda}_i} \geq 0, i \neq 0, s, \quad \text{and } \tilde{\lambda}_i \frac{\partial \tilde{L}}{\partial \tilde{\lambda}_i} = 0, i \neq 0, s,$$

when evaluated at $\{y_s, \theta_s, \tilde{\lambda}^*\}$.

Now suppose $\bar{U}_s < \max_{y,\theta} \mu(\theta) u_s(y)$ subject to $\bar{U}_{s'} \geq \mu(\theta) u_{s'}(y)$, so the constraint $\frac{\mu(\theta)}{\theta} v_s(y) = k$ is binding. That is, $\tilde{\lambda}_0^* > 0$. We conjecture that $\lambda_i^* = \tilde{\lambda}_i^* / \tilde{\lambda}_0^*$ for all $i \neq s$, with $\lambda_0^* = 1$ and $\lambda_s^* = -\tilde{\lambda}_s^* / \tilde{\lambda}_0^*$. Then substituting into the definitions of L and $\tilde{L}$, we find that

$$L(y^*, \theta^*, \lambda^*) = \frac{\tilde{L}(y^*, \theta^*, \tilde{\lambda}^*) - \tilde{\lambda}_s^* \bar{U}_s}{\tilde{\lambda}_0^*} + k.$$

Then it is straightforward to verify that

$$\begin{aligned} \frac{\partial \tilde{L}}{\partial y} &= 0 \Rightarrow \frac{\partial L}{\partial y} = 0, \\ \frac{\partial \tilde{L}}{\partial \theta} &= 0 \Rightarrow \frac{\partial L}{\partial \theta} = 0, \\ \frac{\partial \tilde{L}}{\partial \tilde{\lambda}_i} &\geq 0 \Rightarrow \frac{\partial L}{\partial \lambda_i} \geq 0, i \neq 0, s, \\ \text{and } \tilde{\lambda}_i \frac{\partial \tilde{L}}{\partial \tilde{\lambda}_i} &= 0 \Rightarrow \tilde{\lambda}_i \frac{\partial L}{\partial \lambda_i} = 0, i \neq 0, s. \end{aligned}$$

Also, $\partial L / \partial \lambda_s = (\partial \tilde{L} / \partial \lambda_s - \bar{U}_s) / \tilde{\lambda}_0^* = 0$ since $\partial \tilde{L} / \partial \lambda_s = \mu(\theta_s) u_s(y_s)$. In other words, as long as the constraint on the principal's zero profit condition is binding, $\tilde{\lambda}_0^* > 0$, any solution to problem (P3) also solves problem (P2). ■

It is straightforward to prove that a solution to problem (P3) exists:

Lemma 3 *There exists a vector $\bar{U}$ such that problem* (P3) *is solved for all s. Moreover, if there are multiple solutions, there is one that achieves a higher value than all others.*

Proof. Let $\bar{U}_s^* = \max_{y,\theta} \mu(\theta) u_s(y)$ subject to $\frac{\mu(\theta)}{\theta} v_s(y) = k$. Let $\bar{U}^* = \{\bar{U}_1^*, \ldots, \bar{U}_S^*\}$ and define the mapping $T : [0, \bar{U}^*] \mapsto [0, \bar{U}^*]$ with

$$\begin{aligned} T_s(\bar{U}) = &\max_{y,\theta} \mu(\theta) u_s(y) \\ &\text{s.t. } \frac{\mu(\theta)}{\theta} v_s(y) = k \\ &\text{and } \bar{U}_{s'} \geq \mu(\theta) u_{s'}(y) \text{ for all } s' \neq s. \end{aligned}$$

Since this optimization problem maximizes a continuous function on a compact set, $T_s(\bar{U})$ is defined. T is nondecreasing because higher $\bar{U}_{s'}$ enlarges the constraint set. Given this and the fact that the set $[0, \bar{U}^*]$ is a complete lattice, Tarski's fixed point theorem implies that there exists a fixed point $\bar{U}$ and that the set of fixed points has a largest (and smallest) element. ■

Tarski's Fixed-Point Theorem provides an algorithm for finding the set of equilibria.

Corollary 1 *The highest (lowest) solution to problem* (P3) *can be found by applying the operator T to the highest (lowest) element of the domain, $\bar{U}^*$ ($\mathbf{0}$), and iterating until convergence.*

Now we can establish when an equilibrium exists.

Proposition 2 *Let $\bar{U}_s$ solve problem* (P3) *for all s, with associated allocations $\{y_s, \theta_s\}$. Assume the first constraint in the problem binds for all s. Then $\bar{U}$ and the associated allocations are a competitive search equilibrium.*

Proof. Because the first constraint in problem (P3) binds for all s, $\{y_s, \theta_s\}$ solves problem (P2) and so is a competitive search equilibrium. ■

4 Efficiency

We turn next to our notion of efficiency. We start by defining an incentive compatible and feasible allocation.

Definition 3 *An allocation is a measure λ over the set of incentive-compatible contracts $\mathbb{C}$, with support $\overline{\mathbb{C}}$.*

It is feasible whenever (1) the posted contracts offer the maximal level of expected utility to the agents who are recommended to search for it and no more utility to those who are not; (2) the economy's resource constraint is satisfied; and (3) the recommendations are consistent with the population distribution. Formally:

Definition 4 *Let Λ denote the set of incentive feasible allocations. An allocation is incentive feasible if*

1. *for any $(Y, P, \theta) \in \overline{\mathbb{C}}$, $\bar{U}_s \geq \mu(\theta) u_s(y_s)$ with equality if $p_s > 0$, where*

$$\bar{U}_s \equiv \max_{Y', P', \theta' \in \overline{\mathbb{C}}} \mu(\theta') u_s(y'_s);$$

2. $\int \left(\frac{\mu(\theta)}{\theta} \sum_{s=1}^{S} p_s v_s(y_s) - k \right) d\lambda(Y, P, \theta) = 0$

3. $\frac{1}{\theta} \int p_s d\lambda(Y, P, \theta) \leq \pi_s$, *with equality if* $\bar{U}_s > 0$

An allocation is Pareto optimal if it maximizes the weighted average of expected utilities of agents:

Definition 5 *An allocation is Pareto optimal if it solves* $\max_{\lambda \in \Lambda} \sum_{s=1}^{S} \omega_s \bar{U}_s$ *for some* $\{\omega_s\}$.

To study the Pareto optimal allocations, we use a result similar to the one derived in the equilibrium analysis. The next Lemma states that any Pareto optimal allocation can be represented as a "left separating allocation," that is, an allocation such that each posted contract is designed to attract only one type of agents.[2]

Lemma 4 *Suppose there exists a Pareto optimal allocation* λ*, with associated* $\bar{U}$*, such that some contract* $\mathcal{C} = \{Y, P, \theta\}$ *with* $d\lambda(\mathcal{C}) > 0$ *and either* $p_s, p_{s'} > 0$ *or* $y_s \neq y_{s'}$ *for some* $s \neq s'$*. Let* $\{\{y_s\}, 1_s, \theta\}$ *denote a contract that offers all agents terms of trade* y_s *and attracts only type* s *agents* ($p_s = 1$)*. Then there exists a Pareto optimal allocation* $\tilde{\lambda}$*, associated to the same* $\bar{U}$*, such that* $d\tilde{\lambda}(\mathcal{C}) = 0$*,* $d\tilde{\lambda}(\{y_s\}, 1_s, \theta) = p_s d\lambda(\mathcal{C})$ *for all* s *and* s'*, and* $d\tilde{\lambda}(\mathcal{C}') = d\lambda(\mathcal{C})$ *otherwise.*

Given this Lemma, with a slight abuse of notation, we can represents by $\{y_s, \theta_s\}$ the contract that attracts type s, which implies that $\bar{U}_s = \mu(\theta_s) u_s(y_s)$. Define $\tilde{Y} \equiv \{y_s\}_s$ and $\tilde{\theta} \equiv \{\theta_s\}_s$. Using Lemma 4, we can immediately simplify the characterization of a Pareto optimal allocation:

Proposition 3 *An allocation is Pareto optimal if and only if it can be represented as a pair of* S*-dimensional vectors* $\{\tilde{Y}, \tilde{\theta}\}$ *that solve the following problem:*

$$\max_{\tilde{Y}, \tilde{\theta}} \sum_{s=1}^{S} \omega_s \mu(\theta_s) u_s(y_s) \tag{P4}$$
$$\text{s.t. } \mu(\theta_s) u_s(y_s) \geq \mu(\theta_{s'}) u_s(y_{s'}) \text{ for all } s, s'$$
$$\text{and } \sum_{s=1}^{S} \left(\frac{\mu(\theta_s)}{\theta_s} v_s(y_s) - k \right) \theta_s \pi_s = 0.$$

[2] The proof of this lemma is analogous to the proof of Lemma 1 and therefore omitted.

5 Examples

5.1 Nonexistence

We start by constructing an example in which an equilibrium does not exist. Principals would like to screen agents because some are more productive than others. The problem is that agents have identical preferences and so there is no way to screen them. That is, a principal would always like to announce that he only wants the most productive agents. If he could do so, competition among principals would ensure that more productive agents fare better. But then less productive agents would imitate more productive ones, and so this cannot be an equilibrium. On the other hand, there cannot be an equilibrium where principals do not screen agents, because each principal would request that only the most productive agents seek his contract.

Suppose $S = 2$ and a contract consists only of a transfer from the principal to the agent, $y = t$. Agents' preferences are $u_s(t) = t$ and principals' preferences are $v_s(t) = b_s - t$, with $b_1 < b_2$. Assume $0 \leq t \leq b_2$, so both the principals' and agents' payoffs are continuous and bounded. Since $u_1(t) = u_2(t)$, the constraints in problem (P2) for $s = 1$ requires $\bar{U}_2 \geq \bar{U}_1$, while the constraints in the problem for $s = 2$ require $\bar{U}_1 \geq \bar{U}_2$. Thus any equilibrium must have $\bar{U}_1 = \bar{U}_2$; if the agents have the same preferences, they must get the same level of utility. But the two problems (P2) are not identical. Substituting the constraint $\bar{U}_s = \mu(\theta)t$ into the objective function, the market tightness θ_s must solve $\max_\theta(\mu(\theta)b_s - \bar{U}_s)/\theta$, and both maximized values must equal k. Since $\bar{U}_1 = \bar{U}_2$ and $b_1 \neq b_2$, this is impossible.

It is worth noting what goes wrong in problem (P3) in this case. We look for the best possible equilibrium and so first solve the problem

$$\bar{U}_s = \max_{t,\theta} \mu(\theta)t \text{ s.t. } \frac{\mu(\theta)}{\theta}(b_s - t) = k.$$

That is, we ignore the constraint that the other type of agent must prefer not to seek this contract. Eliminate t using the constraint to get

$$\bar{U}_s = \max_{\theta} \big(\mu(\theta)b_s - \theta k\big).$$

If $b_1 < b_2$, this implies $\bar{U}_1 < \bar{U}_2$. Denote the solution by $\{t_s, \theta_s\}$, $s = 1, 2$. It is easy to verify that the last constraint in problem (P3) for type-2 agents is violated, $\bar{U}_1 < \mu(\theta_2)t_2$.

We therefore must solve the full problem,

$$
\begin{aligned}
\bar{U}_2 = & \max_{t,\theta} \mu(\theta)t \\
\text{s.t. } & \frac{\mu(\theta)}{\theta}(b_2 - t) = k \\
\text{and } & \bar{U}_1 \geq \mu(\theta)t.
\end{aligned}
$$

The Fritz John Lagrangian becomes

$$
\tilde{L}(t,\theta,\tilde{\lambda}) = \tilde{\lambda}_2\mu(\theta)t + \tilde{\lambda}_0\left(\frac{\mu(\theta)}{\theta}(b_2 - t) - k\right) + \tilde{\lambda}_1\big(\bar{U}_1 - \mu(\theta)t\big).
$$

The first order conditions for t and θ are

$$
\mu(\theta_2)(\tilde{\lambda}_2^* - \tilde{\lambda}_1^*) - \frac{\mu(\theta_2)}{\theta_2}\tilde{\lambda}_0^* = 0
$$
$$
\mu'(\theta_2)t_2(\tilde{\lambda}_2^* - \tilde{\lambda}_1^*) + \left(\frac{\mu'(\theta_2)}{\theta_2} - \frac{\mu(\theta_2)}{\theta_2^2}\right)(b_2 - t_2)\tilde{\lambda}_0^* = 0
$$

The solution to these equations has $\tilde{\lambda}_0^* = 0$ and $\tilde{\lambda}_1^* = \tilde{\lambda}_2^* = 1$, with t_2 and θ_2 determined by the solution to the two binding constraints. In particular, this gives $\bar{U}_2 = \bar{U}_1$. The problem for the type-1 agents is unchanged, and so we have found the largest solution to problem (P3). But critically, $\tilde{\lambda}_0^* = 0$ at this solution and so Lemma 2 is inapplicable. In particular, this solution need not, and in this case does not, solve problem (P2).

5.2 Existence of a Least Cost Separating Equilibrium

Our next example introduces a technology for separating the agents.[3] Moving directly to the language of separating contracts, a revelation mechanism now consists of two pieces, $y_s = \{t_s, x_s\}$, where t_s denotes the transfer from the principal to the agent and $x_s \geq 0$ denotes a costly signal that the agent must make. The expected utility of a type-s' agent who searches for a contract $\{t_s, x_s, \theta_s\}$ is

$$
\mu(\theta_s)\left(t_s - \frac{x_s}{a_{s'}}\right),
$$

[3]This example is based on Inderst (2005). We discuss the relationship between his model and ours in more detail below.

where higher values of $a_{s'}$ implies that the action x_s is less costly. The expected profit of a principal who offers this contract is

$$\frac{\mu(\theta)}{\theta}(b_s - t_s) - k,$$

where b_s is the productivity of a type-s agent. Assume $S = 2$ and without loss of generality that type-2 agents are more productive than type-1 agents, $b_2 > b_1$. Also assume that they find the costly signal cheaper to provide, $a_2 > a_1$.[4]

We characterize an equilibrium using problem (P3). Thus we must find numbers $\bar{U}_1$ and $\bar{U}_2$ and contracts $\{t_1, x_1, \theta_1\}$ and $\{t_2, x_2, \theta_2\}$ such that

$$\begin{aligned} \bar{U}_s &= \max_{t,x\geq 0,\theta} \mu(\theta)\left(t - \frac{x}{a_s}\right) \\ \text{subject to } & \frac{\mu(\theta)}{\theta}(b_s - t) = k \\ \text{and } \bar{U}_{s'} &\geq \mu(\theta)\left(t - \frac{x}{a_{s'}}\right) \quad \text{for } s' \neq s. \end{aligned}$$

for $s = 1, 2$. We claim the following result:

Result 1 *There exists a unique competitive search equilibrium, described by*

$$\mu'(\theta_s)b_s = k \textit{ for } s = 1, 2;$$

$$t_s = b_s - \frac{\theta_s}{\mu(\theta_s)}k \textit{ for } s = 1, 2;$$

$$\textit{and } x_1 = 0, \quad x_2 = \frac{a_1}{\mu(\theta_2)}\left[\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right) - \left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right)\right]k.$$

Moreover,

$$\bar{U}_1 = \left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right)k$$

$$\textit{and } \bar{U}_2 = \left[\frac{a_1}{a_2}\left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right) + \left(1 - \frac{a_1}{a_2}\right)\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right)\right]k.$$

We prove this result by finding the best and worst solutions to problem (P3), verifying that they coincide, and verifying that the multiplier on the principal's zero profit condition is positive at the solution, so the solution to problem (P3) also solves the dual problem (P2).

[4] To keep the principals' and agents' payoffs bounded, we can restrict the space of mechanisms to those with $t_s \in [0, b_2]$ and $x_s \in [0, b_2 a_2]$. Any other mechanism would imply negative expected payoffs to either the principal or the agent.

Best Equilibrium. We start by finding the vector $\bar{U}^* = \{\bar{U}_1^*, \bar{U}_2^*\}$ where $\bar{U}_s^*$ solves the problem of maximizing the utility of type-s agents subject to the free entry condition and ignoring the constraint that type-s' agents must be excluded from the market. That is,

$$\bar{U}_s^* = \max_{t,x\geq 0,\theta} \mu(\theta)\left(t - \frac{x}{a_s}\right)$$
$$\text{subject to } \frac{\mu(\theta)}{\theta}(b_s - t) = k.$$

Eliminate t using the constraint to get

$$\bar{U}_s^* = \max_{x\geq 0,\theta} \mu(\theta)\left(b_s - \frac{x}{a_s}\right) - \theta k.$$

The solution sets $x = 0$ and $\theta = \theta_s$, where θ_s solves the necessary and sufficient first order condition

$$\mu'(\theta_s)b_s = k.$$

Since μ is concave, $\theta_1 < \theta_2$. Plugging this back into the original problem, the unconstrained value of agents' solves

$$\bar{U}_s^* = \left(\frac{\mu(\theta_s)}{\mu'(\theta_s)} - \theta_s\right)k.$$

Again using concavity of μ, $\frac{\mu(\theta_s)}{\mu'(\theta_s)} - \theta_s$ is increasing in θ and so $\bar{U}_1^* < \bar{U}_2^*$.

Next we apply the operator T, where

$$T_s(\bar{U}) = \max_{t,x\geq 0,\theta} \mu(\theta)\left(t - \frac{x}{a_s}\right)$$
$$\text{subject to } \frac{\mu(\theta)}{\theta}(b_s - t) = k$$
$$\text{and } \bar{U}_{s'} \geq \mu(\theta)\left(t - \frac{x}{a_{s'}}\right) \text{ for } s' \neq s,$$

starting from $\bar{U} = \bar{U}^*$, to find the best fixed point.

1. First compute $T_1(\bar{U}^*) = \bar{U}_1^*$, since $\bar{U}_2^* > \mu(\theta_1)t_1 = \bar{U}_1^*$. To compute $T_2(\bar{U}^*)$, the same logic implies that the constraint $\bar{U}_1^* \geq \mu(\theta_2)t_2 = \bar{U}_2^*$ is binding. Eliminate t and x from the objective function using the two binding constraints to get

$$\begin{aligned} T_2(\bar{U}^*) &= \max_{\theta}\left[\frac{a_1}{a_2}\bar{U}_1^* + \left(1 - \frac{a_1}{a_2}\right)(\mu(\theta)b_2 - \theta k)\right] \\ &= \left[\frac{a_1}{a_2}\left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right) + \left(1 - \frac{a_1}{a_2}\right)\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right)\right]k, \end{aligned}$$

where still θ_s solves $\mu'(\theta_s)b_s = k$.

2. Now compute $T_1(T(\bar{U}^*)) = T_1(\bar{U}^*) = \bar{U}_1^*$, since $T_2(\bar{U}^*) > \mu(\theta_1)t_1 = T_1(\bar{U}^*)$. That is, the problem for the type-1 agents is still unconstrained by the type-2 agents. On the other hand, since $T_1(\bar{U}^*) = \bar{U}_1^*$, the problem for the type-2 agents is unchanged, so $T_2(T(\bar{U}^*)) = T_2(\bar{U}^*)$. Thus we have found the highest fixed point, the best equilibrium.

Clearly the utility levels coincide with those described in Result 1. It is also straightforward to solve problem (P3) to verify the equilibrium revelation mechanisms.

Worst Equilibrium. Now we start with $\bar{U}_s = 0$ for $s = 1, 2$ to find the worst equilibrium. For low values of $\bar{U}_{s'}$, the inequality constraints in both problems (P3) binds. Use the two binding constraints to eliminate t and x from the problem:

$$\begin{aligned} T_s(\bar{U}) &= \max\left\{0,\ \max_\theta \left[\frac{a_{s'}}{a_s}\bar{U}_{s'} + \left(1 - \frac{a_{s'}}{a_s}\right)(\mu(\theta)b_s - \theta k)\right]\right\} \\ &= \max\left\{0,\ \frac{a_{s'}}{a_s}\bar{U}_{s'} + \left(1 - \frac{a_{s'}}{a_s}\right)\left(\frac{\mu(\theta_s)}{\mu'(\theta_s)} - \theta_s\right)k\right\}. \end{aligned}$$

Note that $T_s(\bar{U}) \geq 0$ because it is always possible to set $\theta_s = \mu(\theta_s) = 0$. We now proceed to apply the operator T.

1. First compute $T(0)$:

$$T_1(0) = 0 \text{ and } T_2(0) = \left(1 - \frac{a_1}{a_2}\right)\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right)k.$$

2. Next compute $T^2(0)$. Since $T_1(0) = 0$, $T_2(T(0)) = T_2(0)$. There are two possibilities for $T_1(T(0))$. First, the constraint $T_2(0) \geq \mu(\theta)(t - x/a_2)$ may be slack in the problem for type-1 agents. In this case, type-1 agents solve the unconstrained problem, giving

$$T_1(T(0)) = \left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right)k.$$

   Alternatively, it is binding, in which case

$$T_1(T(0)) = \left(\frac{a_2}{a_1} - 1\right)\left[\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right) - \left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right)\right]k.$$

3. Now compute $T^3(0)$. Since $T_2(T(0)) = T_2(0)$, $T_1(T^2(0)) = T_1(T(0))$. In the problem

for type-2 agents, the constraint $T_1(T(0)) \geq \mu(\theta)(t - x/a_1)$ will always bind. Thus

$$T_2(T^2(0)) = \frac{a_1}{a_2} T_1(T(0)) + \left(1 - \frac{a_1}{a_2}\right)\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right) k.$$

If in the *previous* step, type-1 agents solved the unconstrained problem, this reduces to

$$T_2(T^2(0)) = \left[\frac{a_1}{a_2}\left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right) + \left(1 - \frac{a_1}{a_2}\right)\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right)\right] k.$$

Note that this means $T_s(T^2(0)) = \bar{U}_s$, $s = 1, 2$, where $\bar{U}_s$ is defined in Result 1. Since the values coincide with the best possible equilibrium, we have found the worst equilibrium, and it coincides with the best. Alternatively, if type-1 agents' problem was constrained in the previous step,

$$T_2(T^2(0)) = \left(1 - \frac{a_1}{a_2}\right)\left[2\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right) - \left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right)\right] k.$$

4. In general, in step n we compute $T^n(0)$. If n is even, $T_2(T^{n-1}(0)) = T_2(T^{n-2}(0))$ since $T_1(T^{n-2}(0)) = T_1(T^{n-3}(0))$. There are always two possibilities for $T_1(T^{n-1}(0))$: the constraint $T_2(T^{n-2}(0)) \geq \mu(\theta)(t - x/a_2)$ may be slack, in which case type-1 agents solve the unconstrained problem, giving

$$T_1(T^{n-1}(0)) = \left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right) k.$$

Or it may be binding, in which case

$$T_1(T^{n-1}(0)) = \frac{n}{2}\left(\frac{a_2}{a_1} - 1\right)\left[\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right) - \left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right)\right] k.$$

5. If n is odd, $T_1(T^{n-1}(0)) = T_1(T^{n-2}(0))$ since $T_2(T^{n-2}(0)) = T_2(T^{n-3}(0))$. If in the previous step type-1 agents solved the unconstrained problem, we obtain $T_s(T^n(0)) = \bar{U}_s$, $s = 1, 2$, where $\bar{U}_s$ is defined in Result 1. Otherwise

$$T_2(T^{n-1}(0)) = \left(1 - \frac{a_1}{a_2}\right)\left[\frac{n+1}{2}\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right) - \frac{n-1}{2}\left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right)\right] k.$$

6. The sequence $T_2(T^{n-2}(0))$, n even, increases linearly in n. This implies that at some finite n, the constraint $T_2(T^{n-2}(0)) \geq \mu(\theta)(t - x/a_2)$ is slack. At this point, we have found the worst equilibrium, and it coincides with the best.

Multiplier on the Zero-Profit Condition. All that remains is to verify that the multiplier on the zero-profit condition is binding. Type-1 agents solve

$$\max_{t,x\geq 0,\theta} \mu(\theta)\left(t - \frac{x}{a_1}\right)$$
$$\text{subject to } \frac{\mu(\theta)}{\theta}(b_1 - t) = k,$$

while type-2 agents solve

$$\max_{t,x\geq 0,\theta} \mu(\theta)\left(t - \frac{x}{a_2}\right)$$
$$\text{subject to } \frac{\mu(\theta)}{\theta}(b_2 - t) = k$$
$$\text{and } \left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right) k \geq \mu(\theta)\left(t - \frac{x}{a_1}\right).$$

It is straight-forward to verify in both cases that the multiplier on the constraint $\frac{\mu(\theta)}{\theta}(b_s - t) = k$ is nonzero, so the value of the problem is higher without this constraint. Thus we have found all the solutions to problem (P3) and proved that any such solution solves problem (P2). This characterizes the set of equilibria.

Alternative Characterization of Equilibrium. An alternative method of finding the fixed point involves less algebra but also does not demonstrate the power of Tarski's fixed point theorem in our environment.

With slight abuse of notation, define the functions $T_s(\bar{U}_{s'})$ with $s' \neq s$ (instead $T_s(\bar{U})$ where $\bar{U} = \{\bar{U}_1, \bar{U}_2\}$. Then $T_s : [0, \bar{U}^*_{s'}] \mapsto [0, \bar{U}^*_s]$. A fixed point is characterized by an intersection of these two functions, such that

$$T_s(T_{s'}(\bar{U}_s)) = \bar{U}_s.$$

The function $T_s(\bar{U}_{s'})$ is defined as follows:

$$T_s(\bar{U}_{s'}) = \max_{t,x\geq 0,\theta} \mu(\theta)\left(t - \frac{x}{a_s}\right)$$
$$\text{subject to } \frac{\mu(\theta)}{\theta}(b_s - t) = k$$
$$\text{and } \bar{U}_{s'} \geq \mu(\theta)\left(t - \frac{x}{a_{s'}}\right).$$

First, notice that if $T_s(\bar{U}_{s'}) = \bar{U}^*_s$, that is, the solution is the unconstrained optimum, then at the optimum $x = 0$. The only condition that needs to be checked to be sure that the unconstrained optimum is feasible is $\bar{U}_{s'} \geq \bar{U}_s^*$. Hence, $T_s(\bar{U}_{s'}) = \bar{U}_s^*$ for any $\bar{U}_{s'} \geq \bar{U}_s^*$. If instead $\bar{U}_{s'} < \bar{U}_s^*$, then the constraint is binding and the problem becomes

$$T_s(\bar{U}_{s'}) = \max_{\theta} \frac{a_{s'}}{a_s}\bar{U}_{s'} + \left(1 - \frac{a_{s'}}{a_s}\right)(\mu(\theta)b_s - \theta k)$$
$$\text{subject to } \bar{U}_{s'} \geq \mu(\theta)b_s - \theta k.$$

Notice that for $s = 1$, the coefficient $1 - a_2/a_1$ is negative and the constraint is binding, that is, $x = 0$, so that $\bar{U}_1 = \bar{U}_2$. For $s = 2$, instead, the coefficient $1 - a_1/a_2$ is positive, so that the solution is

$$\bar{U}_2 = \frac{a_1}{a_2}\bar{U}_1 + \left(1 - \frac{a_1}{a_2}\right)\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right)k,$$

where θ_2 solves

$$\mu'(\theta_2)b_2 = k.$$

Hence the functions are

$$T_1(\bar{U}_2) = \begin{cases} \bar{U}_1^* & \text{if } \bar{U}_2 \geq \bar{U}_1^* \\ \bar{U}_2 & \text{if } \bar{U}_2 < \bar{U}_1^* \end{cases}$$

$$T_2(\bar{U}_1) = \begin{cases} \bar{U}_2^* & \text{if } \bar{U}_1 \geq \bar{U}_2^* \\ \frac{a_1}{a_2}\bar{U}_1 + \left(1 - \frac{a_1}{a_2}\right)\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right)k & \text{if } \bar{U}_1 < \bar{U}_2^* \end{cases}$$

Figure 1 plots these two functions. It is straightforward to see that there exists a unique equilibrium given that $\bar{U}_1^* < \bar{U}_2^*$.

Pareto Improvement. We now propose a feasible allocation that Pareto improves on the Competitive Search Equilibrium. We restrict attention to allocations that treat the two types identically, $t_1 = t_2 = t$, $x_1 = x_2 = x$, and $\theta_1 = \theta_2 = \theta$. Then problem (P4) reduces to

$$\max_{t,x\geq 0,\theta} \mu(\theta)\left(t - \frac{\omega_1 x}{a_1} - \frac{\omega_2 x}{a_2}\right)$$
$$\text{subject to } \mu(\theta)(\pi_1 b_1 + \pi_2 b_2 - t) = \theta k$$

The incentive constraints are automatically satisfied since the contracts are identical. Moreover, it is straightforward to see that $x = 0$, given that it decreases the objective function and does not appear in the constraint. Then eliminate t from the objective using the constraint

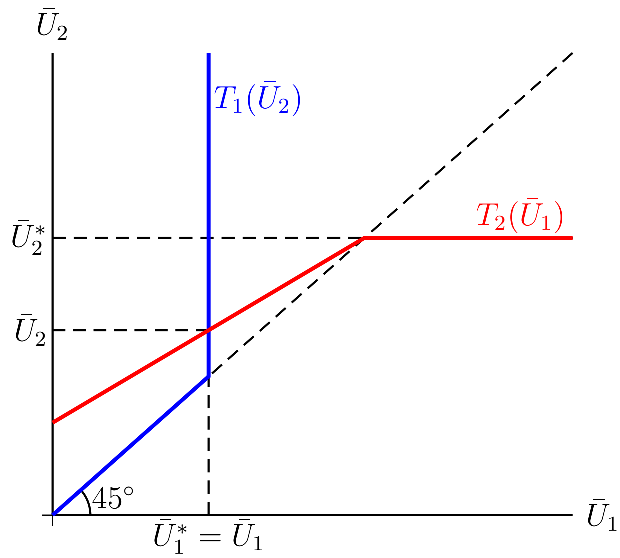

Figure 1: Alternative Characterization of Equilibrium. An equilibrium is at the fixed point $T_1(T_2(\bar{U}_1)) = \bar{U}_1$.

to get

$$\max_{\theta} \big(\mu(\theta)(\pi_1 b_1 + \pi_2 b_2) - \theta k\big).$$

The solution is to set $\theta = \theta^*$ solving $\mu'(\theta^*)(\pi_1 b_1 + \pi_2 b_2) = k$. Since $b_1 < b_2$, $\theta_1 < \theta^* < \theta_2$. Moreover, all agents' expected utility is now

$$\bar{U} = \left(\frac{\mu(\theta^*)}{\mu'(\theta^*)} - \theta^*\right) k.$$

Compare this with the solution in result 1. Since $\theta^* > \theta_1$, trivially $\bar{U} > \bar{U}_1$. On the other hand, $\bar{U} \geq \bar{U}_2$ if and only if

$$\frac{\mu(\theta^*)}{\mu'(\theta^*)} - \theta^* \geq \frac{a_1}{a_2}\left(\frac{\mu(\theta_1)}{\mu'(\theta_1)} - \theta_1\right) + \left(1 - \frac{a_1}{a_2}\right)\left(\frac{\mu(\theta_2)}{\mu'(\theta_2)} - \theta_2\right).$$

This always holds if a_1/a_2 is sufficiently close to 1 (screening is very costly) or if π_1 is sufficiently close to zero (there are a few type-1 agents). The reason is that in equilibrium, principals that want to attract type-2 agents need to screen out type-1 agents. If one principal failed to do so, he would be swamped by them. This may not be socially optimal, however. If there are few type-1 agents or screening is very costly, it is optimal to cross-subsidize type-1 agents and eliminate the need for costly screening.

Relationship to Rothschild and Stiglitz (1976). The structure of this example is similar to the classic insurance model of Rothschild and Stiglitz (1976). Adapting to our environment, they prove that, if there is an equilibrium, it must be the least cost separating equilibrium, where principals that attract type-2 agents use the costly screening technology to keep out type-1 agents. They then argue that there may be a profitable deviation, where principal offers a pooling contract that attracts both types of agents. This occurs exactly in the case when the least separating contract is Pareto inefficient.

Such a deviation is never feasible in our environment. Whereas in Rothschild and Stiglitz (1976), a deviating principal can attract and serve all the agents in the economy, or at least a representative cross-section of agents, this is not the case in our economy. The key is the search friction or capacity constraint. A single principal does not have the ability to serve all the agents who would potentially be attracted to a contract. Instead they are rationed thorough the endogenous movement in market tightness θ. The key to whether such a deviation is profitable then is understanding which agents are most willing to accept a decline in market tightness. In this model, type-2 agents will quickly give up on the pooling contract if it is too crowded with type-1 agents. Type-1 agents, who have a lower outside option, $\bar{U}_1 < \bar{U}_2$, are more persistent. A principal who tries to offer the pooling contract proposed by Rothschild and Stiglitz (1976) will find himself with a long queue of type-1 agents, the worst possible outcome.

The exact environment studied by Rothschild and Stiglitz (1976) is somewhat different than this model. In their framework, risk-neutral principals offer insurance contracts to risk-averse agents with heterogeneous probabilities of loss. Our general setting is rich enough to encompass that model and one can verify that a separating equilibrium always exists, but may be Pareto inefficient. The reasons are exactly the same as in the simple model we outline here.

We borrowed this example from Inderst (2005). He studies a dynamic model with random meetings between principals and agents. There is an exogenous inflow of new principals and agents into the marketplace, and when principals and agents agree on a contract, they exit the market. He finds that in steady state, there will always be sufficiently many type-1 (undesirable) agents that principals will offer screening contracts. That is, the endogenous composition of the pool of agents circumvents the Rothschild and Stiglitz (1976) nonexistence result. We treat the composition of the entire pool of agents as exogenous; however, through her contract, the composition of the pool of agents attracted to a particular principal is endogenous. A principal that tries to offer a pooling contract finds the composition of the market turning against her.

5.3 Distortions in Market Tightness

A peculiar feature of the previous example is that market tightness is not distorted by the need to screen. That is, θ_s is always at the constrained efficient level, maximizing $\mu(\theta)b_s - \theta k$. In the next example, we study how principals can use market tightness to screen out the undesired type. To stress this point, we assume $\mu(\theta) = \min\{\theta, 1\}$. Thus each principal can match with only one agent and vice versa, but there are no other coordination or search frictions.

We imagine that each agent is endowed with one indivisible apple, but type-1 agents have a bad apple while type-2 agents have a good apple. Each principal is endowed with one indivisible banana. A trading mechanism for type-s agents is a pair $\{\alpha_s, \beta_s\}$, where α_s is the probability that the agent gives the principal her apple and β_s is the probability that the principal gives the agent her banana. The expected utility of a type-s' agent seeking the contract for a type-s agent is

$$\mu(\theta_s)\big(\beta_s b - \alpha_s a^A_{s'}\big),$$

where $b > 0$ is the utility she derives from consuming a banana and $a^A_s > 0$ is the value that a type-s agent places on consuming her apple. The expected utility of a principal offering a type-s contract is

$$\frac{\mu(\theta_s)}{\theta_s}\big(\alpha_s a^P_s - \beta_s b\big) - k.$$

Here $a^P_s > 0$ is the value that a principal places on consuming a type-s apple, while we normalize her value of a banana to the same b. Assume $S = 2$ and without loss of generality that $a^P_2 > a^P_1$. Also assume $a^P_s > a^A_s + k$ for $s = 1, 2$, so there are gains from trade for both types of apples, and $a^P_s < b + k$, so principals are unwilling to post a contract that gives up their entire banana in return for an apple.

We look for a solution to problem (P3), which we state here as

$$\bar{U}_s = \max_{\alpha\in[0,1],\beta\in[0,1],\theta} \mu(\theta)\big(\beta b - \alpha a^A_s\big)$$

$$\text{subject to } \frac{\mu(\theta)}{\theta}\big(\alpha a^P_s - \beta b\big) = k$$

$$\text{and } \bar{U}_{s'} \geq \mu(\theta)\big(\beta b - \alpha a^A_{s'}\big) \text{ for } s' \neq s.$$

Eliminate β using the free-entry condition to write this as

$$\bar{U}_s = \max_{\alpha\in[0,1],\theta} \mu(\theta)\alpha\big(a_s^P - a_s^A\big) - \theta k$$
$$\text{subject to } \alpha a_s^P - \frac{\theta k}{\mu(\theta)} \in [0,b]$$
$$\text{and } \bar{U}_{s'} \geq \mu(\theta)\alpha\big(a_s^P - a_{s'}^A\big) - \theta k \text{ for } s' \neq s.$$

We again look for the best equilibrium. We start by solving the unconstrained problem for type-s agents:

$$\bar{U}_s^* = \max_{\alpha\in[0,1],\theta} \mu(\theta)\alpha\big(a_s^P - a_s^A\big) - \theta k$$
$$\text{subject to } \alpha a_s^P - \frac{\theta k}{\mu(\theta)} \in [0,b].$$

Since $\mu(\theta) = \min\{\theta, 1\}$, the solution is to set $\alpha = \theta = 1$, delivering $\bar{U}_s^* = a_s^P - a_s^A - k > 0$. Next we apply the operator T, where

$$T_s(\bar{U}) = \max_{\alpha\in[0,1],\theta} \mu(\theta)\alpha\big(a_s^P - a_s^A\big) - \theta k$$
$$\text{subject to } \alpha a_s^P - \frac{\theta k}{\mu(\theta)} \in [0,b]$$
$$\text{and } \bar{U}_{s'} \geq \mu(\theta)\alpha\big(a_s^P - a_{s'}^A\big) - \theta k \text{ for } s' \neq s.$$

Starting from $\bar{U} = \bar{U}^*$, we guess that the constraint for type-1 agents is slack while the one for type-2 agents is binding. Thus

$$T_1(\bar{U}^*) = \bar{U}_1^*$$

and

$$T_2(\bar{U}) = \max_{\alpha\in[0,1],\theta} \mu(\theta)\alpha\big(a_2^P - a_2^A\big) - \theta k$$
$$\text{subject to } \alpha a_2^P - \frac{\theta k}{\mu(\theta)} \in [0,b]$$
$$\text{and } a_1^P - a_1^A - k = \mu(\theta)\alpha\big(a_2^P - a_1^A\big) - \theta k.$$

Eliminate α:

$$\begin{aligned} T_2(\bar{U}) = &\max_\theta \frac{a_1^P - a_1^A - (1-\theta)k}{a_2^P - a_1^A}\left(a_2^P - a_2^A\right) - \theta k \\ \text{subject to } &\frac{a_1^P - a_1^A - (1-\theta)k}{\mu(\theta)(a_2^P - a_1^A)} a_2^P - \frac{\theta k}{\mu(\theta)} \in [0, b] \\ \text{and } &\frac{a_1^P - a_1^A - (1-\theta)k}{\mu(\theta)(a_2^P - a_1^A)} \in [0, 1] \end{aligned}$$

Temporarily ignore the constraints, which ensure that α and β are proper probabilities. Since $a_2^P > a_1^A$ by assumption, the objective function is increasing in θ if and only $a_1^A > a_2^A$. This gives two cases, depending on whether this inequality holds.

If $a_1^A < a_2^A$, the objective function is decreasing in θ. Set θ equal to the smallest value consistent with the two constraints. Here that is

$$\theta_2 = \frac{a_1^P - a_1^A - k}{a_2^P - a_1^A - k} < 1,$$

which implies $\alpha_2 = 1$, so the second constraint binds, and

$$\bar{U}_2 = \frac{(a_2^P - a_2^A - k)(a_1^P - a_1^A - k)}{a_2^P - a_1^A - k}.$$

It is straightforward to verify that the first constraint is satisfied. In addition, one can verify that the principals' zero profit condition binds and so the solution to problem (P3) also solves (P2). Principals create too few contracts, so some type-2 agents fail to meet a principal. Since type-2 agents hold a better apple than do type-1 agents, they are more willing to accept this low meeting probability. In return, they get more bananas for their apples when they succeed in meeting a principal. Note that the obvious alternative, setting $\theta_s = 1$ but rationing though the probability of exchange, $\alpha_s < 1$, is more costly because it involves creating more contracts at cost k per contract.

Conversely, if $a_1^A > a_2^A$, the objective function is increasing in θ. Now set θ equal to the largest value consistent with the two constraints,

$$\theta_2 = 1 + \frac{a_2^P - a_1^P}{k} > 1,$$

which again implies $\alpha_2 = 1$. The value of the objective function is

$$\bar{U}_2 = a_1^P - a_2^A - k$$

Again, it is straightforward to verify that the first constraint is satisfied. However, in this case principals' zero profit condition is slack and so Lemma 2 is inapplicable. The solution to problem (P3) may not solve the primal problem (P2).

Indeed, this is the case. Suppose there were an equilibrium with $\bar{U}_1 = a_1^P - a_1^A - k$ and $\bar{U}_2 = a_1^P - a_2^A - k$. Write problem (P2) for the type-2 contracts:

$$\begin{aligned} k = \max_{\alpha\in[0,1],\beta\in[0,1],\theta} & \frac{\mu(\theta)}{\theta}(\alpha a_2^P - \beta b) \\ \text{subject to } & \bar{U}_2 = \mu(\theta)(\beta b - \alpha a_2^A) \\ \text{and } & \bar{U}_1 \geq \mu(\theta)(\beta b - \alpha a_1^A). \end{aligned}$$

It is feasible to set $\theta = \alpha = 1$ and $\beta b = a_1^P - k$, since this implies both constraints bind. But then the value of the objective function is $a_2^P - a_1^P + k > k$, a contradiction.

When $a_1^A > a_2^A$, there is no equilibrium because the screening technology is useless. Principals would prefer to obtain type-2 apples, while type-2 agents are more willing to part with their apple. Since the only means of screening is through probabilistic trade, type-1 agents are always willing to incur any screening costs that type-2 agents find palatable.

Pareto Improvement. It is again possible to attain a Pareto improvement. To show this, we find a feasible allocation that Pareto dominates the equilibrium. We focus on allocations that offer the same contract to different types, with $\alpha_1 = \alpha_2 = \theta_1 = \theta_2 = 1$. Then problem (P4) reduces to

$$\begin{aligned} & \max_{\beta\in[0,1]} \left(\beta b - \omega_1 a_1^A - \omega_2 a_2^A\right) \\ \text{s.t. } & \pi_1 a_1^P + \pi_2 a_2^P - \beta b - k = 0. \end{aligned}$$

Using the constraint to eliminate βb, the expected utility of a type-s agent is

$$\bar{U}_s = \pi_1 a_1^P + \pi_2 a_2^P - a_s^A - k.$$

Since $\bar{U}_s^* = a_s^P - a_s^A - k$ and $a_1^P < a_2^P$, it is immediate that type-1 agents are better off. Type-2 agents are better off if

$$\pi_1 a_1^P + \pi_2 a_2^P - a_2^A - k > \frac{(a_2^P - a_2^A - k)(a_1^P - a_1^A - k)}{a_2^P - a_1^A - k},$$

which holds for all π_s whenever $a_1^A \leq a_2^A$. Thus whenever an equilibrium exists, it would be better for both types of agents to ensure that everyone trades at the unconditional fair price.

References

Acemoglu, Daron and Robert Shimer, "Efficient Unemployment Insurance," *Journal of Political Economy*, 1999, *107* (5), 893–928.

Faig, Miguel and Belen Jerez, "A Theory of Commerce," *Journal of Economic Theory*, 2005, *122* (1), 60–99.

Guerrieri, Veronica, "Efficiency of Competitive Search under Asymmetric Information," *Journal of Political Economy*, 2008. forthcoming.

Inderst, Roman, "Matching Markets with Adverse Sselection," *Journal of Economic Theory*, 2005, *121* (2), 145–166.

Moen, Espen R., "Competitive Search Equilibrium," *Journal of Political Economy*, 1997, *105* (2), 385–411.

_ **and Åsa Rosén**, "Incentives in Competitive Search Equilibrium and Wage Rigidity," 2006. mimeo.

Montgomery, James D., "Equilibrium Wage Dispersion and Interindustry Wage Differentials," *Quarterly Journal of Economics*, 1991, *106* (1).

Peters, Michael, "Ex Ante Price Offers in Matching Games: Non-Steady States," *Econometrica*, 1991, *59* (5), 1425–1454.

Rothschild, Michael and Joseph Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 1976, *90* (4), 629–649.

Shimer, Robert, "Contracts in a Frictional Labor Market," 1996. MIT mimeo.

Simon, Carl P. and Lawrence E. Blume, *Mathematics for Economists*, New York: Norton, 1994.