

Competitive Equilibria in Markets with Adverse Selection

Let us consider markets, such as insurance markets, where each consumer can buy a contract from one firm, which would establish a contractual relationship between the consumer and the firm. In this relationship, the consumer's expected benefit and the firm's expected cost of fulfilling the terms of any contract may depend on the consumer's type, which is privately known only by the consumer. Thus, in selling any particular contract, a firm must worry about the possibility that it is selling to consumers who might generate particularly high costs for the firm under this contract; this is the problem of *adverse selection*.

We assume that there are multiple firms which can sell any of these contracts, and any one of these firms could serve the entire population of consumers. Thus, in a competitive equilibrium, prices should be such that firms expect zero profits for every contract, when problems of adverse selection are taken appropriately into account. Our goal is to understand what competitive equilibria could look like in such a market.

Michael Rothschild and Joseph Stiglitz (*QJE*, 1976) considered one natural definition of competitive equilibrium which they found might not exist for some simple markets. But if economists cannot even identify any competitive equilibria for simple models of markets with adverse selection, then economic analysis cannot contribute to policy debates about whether insurance markets would function better with more competition or with more regulation. So we need a more refined understanding of perfect competition that would exclude at least some of the competitive deviations which Rothschild and Stiglitz used to eliminate all candidates for equilibrium in the examples where Rothschild and Stiglitz found nonexistence.

Here we follow Eduardo Azevedo and Daniel Gottlieb (*Econometrica*, 2017) and develop a simplified version of their concept of competitive equilibrium for markets with adverse selection, which they proved to exist for a wide class of market models. This approach differs from Rothschild and Stiglitz in not considering competitive deviations in which a firm substantially cuts its price for one kind of contract in order to attract profitable demand from low-cost consumers who would not have otherwise considered buying this contract. For a firm to attract demand from consumers who had not been planning to buy this contract at the anticipated market price, the price cut would have to be advertised broadly to the general population of consumers, and such a broad advertising campaign could easily provoke counter-responses from other firms.

Economists generally assume that in a perfectly competitive market, consumers plan their consumption optimally with respect to the given market prices and then can readily find many alternative sources of supply at these prices. In a perfectly competitive market, each firm gets only a small share of the overall demand for any positively traded contract, but this share would be very elastic to any deviation from the equilibrium price. The ability of a firm to substantially increase its share of existing demand by offering to sell a contract for even slightly less than the market price is sufficient to imply that firms must expect to just break even on all contracts in a competitive equilibrium. Thus, competitive equilibria here are characterized by an understanding of which types of consumers would buy each contract at its prevailing market prices such that no firm could expect to gain by a small price cut to increase its own share of the demand from these consumers.

A basic framework for analyzing markets with adverse selection

To model such a market with adverse selection, let Y denote the set of possible contracts, and let T denote the set of possible types of consumers. For simplicity, we may sometimes assume that Y and T are nonempty finite sets. (When they are infinite sets, some sums below may need to be rewritten as integrals.) A given utility function $U: \mathbb{R} \times Y \times T \rightarrow \mathbb{R}$ specifies the utility $U(p, x, t)$ that any consumer of type t would get from buying contract x at price p . A given cost function $C: Y \times T \rightarrow \mathbb{R}$ specifies the expected cost $C(x, t)$ for a firm to fulfill a contract x for a consumer of type t . A given probability distribution μ in $\Delta(T)$ specifies the fraction $\mu(t) > 0$ of consumers who have each type t in the general population. Each consumer must buy exactly one contract. (The set of possible contracts Y may include a "no-trade" option $x=0$ which would have zero cost $C(0, t)=0$ regardless of the consumer's type t .)

In a Rothschild-Stiglitz model of insurance markets, a consumer's type $t \in [0, 1]$ would denote his probability of suffering some loss $\ell > 0$ from some given initial wealth W , and the contract parameter $x \in [0, 1]$ would denote the fraction of this loss to be covered by an insurance policy. Then, given some concave increasing utility function $u(\bullet)$ for monetary wealth, we would get

$$U(p, x, t) = t u(W - p - (1-x)\ell) + (1-t) u(W - p).$$

In general, we assume here that, for each $x \in Y$ and each $t \in T$, $U(p, x, t)$ is strictly decreasing and continuous in the price p . Also, for each $x \in Y$, $t \in T$, $p \in \mathbb{R}$, and $y \in Y$, we assume that there exists some price $\varphi(p, x, t, y)$ such that $U(p, x, t) = U(\varphi(p, x, t, y), y, t)$. (This says that money is important enough for a price adjustment to change any consumer's preference over any pair of contracts.) So $\varphi(p, x, t, y)$ denotes the price of y which would make a type- t consumer indifferent between buying y and buying x at price p .

We consider here a simplified version of Azevedo and Gottlieb's (2017) definition of competitive equilibrium for markets with adverse selection. We define a competitive equilibrium to be a pair (q, γ) such that $q = (q(x))_{x \in Y}$ is a price vector in \mathbb{R}^Y , $\gamma = (\gamma(x|t))_{x \in Y, t \in T}$ is an allocation vector in $\Delta(Y)^T$, and the following conditions are satisfied:

- [0] $\sum_{y \in Y} \gamma(y|t) = 1$ and $\gamma(x|t) \geq 0, \forall x \in Y, \forall t \in T$;
- [1] $\sum_{x \in Y} \gamma(x|t) U(q(x), x, t) = \max_{x \in Y} U(q(x), x, t), \forall t \in T$;
- [2] $\sum_{t \in T} \mu(t) \gamma(x|t) (C(x, t) - q(x)) = 0, \forall x \in Y$; and
- [3] $\forall y \in Y, \exists t \in T$ such that $U(q(y), y, t) = \max_{x \in Y} U(q(x), x, t)$ and $q(y) \leq C(y, t)$.

Condition [0] must be satisfied because each number $\gamma(x|t)$ denotes the fraction of type- t consumers who choose contract x . Condition [1] is an optimality condition for consumers, saying that consumers of each type only choose contracts that maximize their utility at the given q prices. Condition [2] is a zero-profit condition for firms, saying that firms expect to break even on any contract that attracts a positive fraction of the consumers. In this case, when $\sum_{s \in T} \mu(s) \gamma(x|s) > 0$, condition [2] implies that $q(x)$ equals the average cost of all consumers who choose contract x

$$q(x) = \sum_{t \in T} C(x, t) \mu(t) \gamma(x|t) / \sum_{s \in T} \mu(s) \gamma(x|s).$$

Condition [3] says that, for each contract, there is at least one type which is willing to choose this contract but would not be profitable for the competitive firms at the given price. Condition [3] is actually implied by condition [2] for any contract that has positive demand in the equilibrium, but requiring condition [3] for all contracts adds a further restriction on the pricing of contracts

that have zero demand under γ . Without this restriction, we could get equilibria that satisfy [2] trivially for any one contract x simply by setting its price $q(x)$ so high that all $\gamma(x|t)=0$. With condition [3], a real possibility of unprofitable sales can explain what deters firms from trying to increase demand for a contract by shading its price slightly.

When all costs $C(x,t)$ are nonnegative, any positively traded contract must have a nonnegative price in equilibrium; but the definition of competitive equilibrium here allows the possibility that some untraded contracts might have negative prices, as if there were some small fund for subsidizing sales of the contract. Of course, a contract that has zero demand at a negative price would have the same zero demand at the higher price of zero. The point of imputing a negative price here is only to identify which types would be willing to buy the contract with the least subsidy, so that we can verify that the imputed price is not greater than the costs of serving these types.¹

We now state and prove a general existence theorem for equilibria of finite markets with adverse selection. To simplify some equations below, let us introduce the notation:

$$\bar{U}(q,t) = \max_{x \in Y} U(q(x),x,t).$$

Fact. When Y and T are finite sets and U satisfies the assumptions that are listed above, a competitive equilibrium must exist.

Proof. For each x in Y , let $\bar{c}(x) = \max_{t \in T} C(x,t)$, and $\underline{c}(x) = \min_{t \in T} C(x,t)$.

Then let $c_0(x) = \min_{t \in T} \min_{y \in Y} \varphi(\underline{c}(y),y,t,x) - 1$.

Here we have $c_0(x) < \underline{c}(x)$, because $\varphi(\underline{c}(x),x,t,x) = \underline{c}(x)$, and we also have $\underline{c}(x) \leq \bar{c}(x)$. Now, for any ε such that $0 < \varepsilon < 1$, consider a modified market in which the fraction of each type t in T is $(1-\varepsilon)\mu(t)$ and, on for each contract x , a fraction $\varepsilon/\#Y$ of the consumers are a new artificial type that only buys contract x and has cost $c_0(x)$. (Here $\#Y$ is the number of contracts in the finite set Y .)

Now for any price vector q in $\times_{x \in Y} [c_0(x),\bar{c}(x)]$ and any allocation vector γ in $\Delta(Y)^T$ such that, consider a mapping that selects a new price vector q' and a new allocation vector γ' as follows.

For each contract x , $q'(x)$ is the average cost

$$q'(x) = (c_0(x)\varepsilon/\#Y + \sum_{t \in T} C(x,t)\mu(t)\gamma(x|t)) / (\varepsilon/\#Y + \sum_{t \in T} \mu(t)\gamma(x|t)).$$

For each type t , $\gamma'(\bullet|t)$ can be any probability distribution over Y such that

$$\{x \in Y \mid \gamma'(x|t) > 0\} \subseteq \operatorname{argmax}_{x \in Y} U(q(x),x,t), \quad \forall t \in T.$$

By the Kakutani fixed-point theorem, we can find a pair $(q^\varepsilon, \gamma^\varepsilon)$ in $(\times_{x \in Y} [c_0(x),\bar{c}(x)]) \times \Delta(Y)^T$ that maps to itself under this mapping. This $(q^\varepsilon, \gamma^\varepsilon)$ will satisfy conditions [0] and [1] and also a perturbed version of

¹ The definition of competitive equilibrium here is a simplified version of the equilibrium concept developed by Azevedo and Gottlieb [2017]. Their concept differs from the simple competitive equilibrium concept here in that they require that the prices of untraded contracts should be evaluated as limits of equilibrium prices from a sequence of perturbed models in which vanishingly small populations of artificial consumers are introduced to guarantee some positive low-cost demand for every contract. (Such a construction is used in the proof of equilibrium existence above.) This perturbational condition is analogous to conditions that are applied in perfect and proper refinements of Nash equilibrium. These perturbations are constructed so that the limiting equilibrium price of each contract is nonnegative, and so the type that would be willing to buy the contract with the least subsidy might not be identified when the price is zero, but the costs for all types are assumed to be nonnegative. Examples can be constructed (with infinitely many types, including pairs of types that have identical utility functions but different costs) such that the perfect competitive equilibrium concept of Azevedo and Gottlieb helps to exclude some counter-intuitive equilibria that would be admitted by the simple competitive equilibrium concept developed here.

the zero-profit condition [2] with the new artificial types included as an ε fraction of the population (which ensure that every contract gets strictly positive demand).

Now consider a sequence of numbers ε that converge to 0, so that the artificial types become an infinitesimal fraction of the population as we go to the limit. By compactness of the domain $(\times_{x \in Y} [c_0(x), \bar{c}(x)]) \times \Delta(Y)^T$, there exists a subsequence such that all the $q^\varepsilon(x)$ and $\gamma^\varepsilon(x|t)$ converge to some limits $q^*(x)$ and $\gamma^*(x|t)$ that are also in this domain. We can now show that this (q^*, γ^*) is a competitive equilibrium.

The fact that conditions [0] and [1] and an approximate version of [2] are satisfied for each $(q^\varepsilon, \gamma^\varepsilon)$ implies, by continuity, that these conditions are satisfied for the limit (q^*, γ^*) . Furthermore, the limit of $\varepsilon \rightarrow 0$ implies that every contract x that has positive γ^* -demand must have $q^*(x)$ that is a weighted average of $C(x, t)$ costs and so satisfies $\underline{c}(x) \leq q^*(x) \leq \bar{c}(x)$. Each type t must be generating positive demand $\gamma(x|t) > 0$ for at least one of its optimal contracts x , for which we then get that t 's optimal utility $\bar{U}(q^*, t)$ must satisfy $\bar{U}(q^*, t) = U(q^*(x), x, t) \leq U(\underline{c}(x), x, t)$. (Utility is decreasing in the price.)

Now let us show that condition [3] is satisfied for any contract y . Since $c_0(y) < \underline{c}(y)$, the price $q^*(y)$ could not equal $c_0(y)$ unless demand for y was zero. But $c_0(y)$ was defined so that any type would strictly prefer y at price $c_0(y)$ over any other contract x priced at $\underline{c}(x)$. Thus, for every contract y , we must have $q^*(y) > c_0(y)$, which implies $q^\varepsilon(y) > c_0(y)$ for all sufficiently small ε , which in turn implies that $\gamma^\varepsilon(y|t) > 0$ for at least one real type t in T . (The demand $\gamma^\varepsilon(y|t)$ would be small, of order ε , but it must be strictly positive.) As there are only finitely many types, we can choose a subsequence (if necessary) so that there is some type t in T such that $\gamma^\varepsilon(y|t) > 0$ for all ε , and so we get $U(q^\varepsilon(y), y, t) = \bar{U}(q^\varepsilon, t)$ for all ε . Among such types t , we can pick the one with the highest cost in contract y , and then we must also get $q^\varepsilon(y) \leq C(y, t)$ for all ε , because the price $q^\varepsilon(y)$ is a weighted average of the real types that choose y in γ^ε and the lower-cost artificial type for y . So in the limit, the price $q^*(y)$ satisfies condition [3], $U(q^*(y), y, t) = \bar{U}(q^*, t)$ and $q^*(y) \leq C(y, t)$, with this type t .

Thus (q^*, γ^*) is a competitive equilibrium. **QED**

Separating equilibria of markets with two risk-types and a convex interval of contracts

For markets where the set of possible contracts Y is infinite, one can still prove an existence of competitive equilibria with some additional continuity assumptions which are developed by Azevedo and Gottlieb (2017). The basic idea of the proof is to consider the limits of competitive equilibria that we would get for an increasing sequence of finite subsets of Y which in the limit become dense in the whole set.

Now let us consider a class of models that include Rothschild-Stiglitz insurance markets with two types. Let $T = \{L, H\}$ where L is the low type and H is the high type. The type may be interpreted as the probability of suffering a loss, with $0 < L < H < 1$. The set of possible contracts Y will be the interval $[0, 1]$ (or some subset of this interval), where any x in $[0, 1]$ denotes the fraction of the insured loss that will be covered by the insurance company. In addition to the previous assumptions about U , we now add several further assumptions about the U and C functions in our model, all of which will be satisfied by a Rothschild-Stiglitz insurance model:

$C(x, t)$ is continuous in x and $U(p, x, t)$ is continuous in x and p , $\forall t \in T$;

$C(0, L) = C(0, H) = 0$, but $0 < C(x, L) < C(x, H) \forall x > 0$;

$U(C(x, t), x, t)$ is strictly increasing in x , $\forall t \in T$;

if $y > x$ then $\varphi(p, x, L, y) < \varphi(p, x, H, y)$, but if $y < x$ then $\varphi(p, x, L, y) > \varphi(p, x, H, y)$.

The second assumption says that the high types are more expensive to serve. The third assumption says that any type of consumer would prefer full insurance at a price that is actuarially fair for this type. The fourth assumption is a "single-crossing" condition. It says that, when consumers are asked about their willingness to switch from a contract x at a price p to some other contract y , if y is greater than x then type H would be more willing to switch than type L, in the sense that type H would accept the increased coverage y for a higher price than type L would accept; but if $y < x$ then the type L would be more willing to switch than type H. If U is continuously differentiable in both p and x , then the single-crossing property is implied by a condition that, when we normalize the marginal utility of money ($-p$) across types, type H has greater marginal utility for increasing coverage x .

$$(\partial U(p,x,L)/\partial x)/(-\partial U(p,x,L)/\partial p) < (\partial U(p,x,H)/\partial x)/(-\partial U(p,x,H)/\partial p).$$

Fact. When $Y=[0,1]$, $T=\{L,H\}$, and U and C satisfy the above assumptions, the unique competitive equilibrium is the best separating plan, which is defined as follows. Type H chooses the contract $x_H=1$ which has price $q(1)=C(1,H)$. Type L chooses a contract x_L at the price $q(x_L) = C(x_L,L)$, where x_L is the solution to the equation $U(C(x_L,L),x_L,H) = U(C(1,H),1,H)$. Any other contract y has the price $q(y) = \varphi(C(x_L,L),x_L,L,y)$ if $y < x_L$, and $q(y) = \varphi(C(1,H),1,H,y)$ if $y > x_L$.

Proof. Every contract x in the interval $[0,1]$ must be priced so that at least one type is willing to buy it, so that either $U(q(x),x,L) = \bar{U}(q,L)$ or $U(q(x),x,L) = \bar{U}(q,H)$. The pricing function $q(x)$ cannot be discontinuous because, at any point of discontinuity, the type that is (supposedly) willing to choose contracts that converge to the highest price limit would actually strictly prefer to switch to a nearby contract that has a price close to the lowest price limit. Thus, the set of contracts that any type is willing to choose at the q prices is a closed set, and so the connected interval $Y=[0,1]$ must include some contract y^* that both types are willing to choose. That is, we have some y^* such that $U(q(y^*),y^*,L) = \bar{U}(q,L)$ and $U(q(y^*),y^*,H) = \bar{U}(q,H)$.

By the single-crossing property, this y^* must be unique, and only type L is willing to choose any contract $x < y^*$. (If $x < y^*$ satisfied $U(q(x),x,H) = \bar{U}(q,H) = U(q(y^*),y^*,H)$, then type L would strictly prefer x at price $q(x)$ over y^* at price $q(y^*)$, contradicting the fact that y^* is optimal for both types.) Similarly, if $x > y^*$, then only type H is willing to choose the contract x at price $q(x)$.

So for any $\varepsilon > 0$, condition [3] requires $U(q(y^*-\varepsilon),y^*-\varepsilon,L) = \bar{U}(q,L)$ and $q(y^*-\varepsilon) \leq C(y^*-\varepsilon,L)$. Taking the limit as $\varepsilon \rightarrow 0$, we must have $q(y^*) \leq C(y^*,L)$. The contracts that L actually chooses must be in the interval $[0,y^*]$. If L actually had positive demand for some contract $x < y^*$, then the zero-profit condition for firms would require that contract to be priced at $q(x)=C(x,L)$, but $U(C(x,L),x,L)$ is strictly increasing in x , and so we would get the contradiction

$$\bar{U}(q,L) = U(q(x),x,L) = U(C(x,L),x,L) < U(C(y^*,L),y^*,L) \leq U(q(y^*),y^*,L) = \bar{U}(q,L).$$

Thus type L cannot be actually buying any contract less than y^* , which is the highest contract that L is willing to choose. We found that $q(y^*) \leq C(y^*,L)$, but for firms to break even in selling y^* , its price cannot be strictly less than the low cost $C(y^*,L)$, and so we must have $q(y^*) = C(y^*,L)$, and only the type-L consumers are actually buying this contract. That is, y^* is the contract that was called x_L in the statement of the Fact above.

Then type H consumers must be actually buying some strictly higher contract $z > y^*$ which must satisfy the break-even price of $q(z)=C(z,H)$. But $U(C(z,H),z,H)$ is strictly increasing in z , and we know that the greatest contract 1, which only H is willing to choose, must in equilibrium satisfy $q(1) \leq C(1,H)$. Thus, if type H cannot be buying any contract $z < H$, because it would yield the contradiction $\bar{U}(q,H) = U(C(z,H),z,H) < U(C(1,H),1,H) \leq U(q(1),z,H) = \bar{U}(q,H)$. Thus, type H is buying the maximal contract 1 at the price $C(1,H)$. **QED**

An example where competitive equilibria are inefficient

In insurance markets that fit this two-type model, the high-risk types get full insurance priced for their higher risks, and the low-risk types get substantially incomplete insurance although it is fairly priced for their lower risks. The low-risk types must accept less insurance coverage as a costly signal to distinguish themselves from high-risk types. But notice (in the *Fact* above) that the contract that each type gets and the price that it pays in these models actually do not depend on the fraction of each type in the population. If most consumers are the high-risk type, then it is not surprising that low-risk types would choose to accept less insurance coverage as a signaling cost to distinguish themselves from the high-risk types, because the low types would have to pay much more for insurance if they had to be pooled with all those high types. However, when the fraction of high-risk types in the overall population is small, then the average cost of providing full insurance to everybody in the population may be only slightly greater than the cost of providing it to the low-risk types separately. In that case, the separating equilibrium will be Pareto-inefficient, as all types would strictly prefer to buy full (or almost-full) insurance at a price that fairly covers the pooled costs of the whole population.

In such a case, a firm might hope to earn positive profits by offering full insurance at a price that is slightly above the average cost for the entire population, if the firm could actually attract the entire population without stimulating a competitive response by other firms. But such a pooling offer would not be an equilibrium, because other firms could offer slightly less insurance coverage at a slightly lower price that would profitably attract only the low-risk types, as long as the high-risk types have the option of accepting the first firm's "pooling" offer. The Pareto-efficient pooling contract could be sustained as a competitive equilibrium only with some government intervention to prevent firms from offering contracts with slightly less coverage.

For example consider an insurance market where each consumer has an independent risk of some insurable loss of size $\ell=30$, and each consumer has constant risk tolerance $\tau=10$. In the population, 75% of all consumers are low-risk types for whom the probability of this loss is 0.2, but 25% are high-risk types for whom the probability of this loss is 0.6. So the cost of an insurance contract that covers a fraction x of the loss for each type would be $C(x,L)=0.2x\ell=6x$ and $C(x,H)=0.6x\ell=18x$, while the pooled cost would be $0.75(0.2x\ell)+0.25(0.6x\ell) = 0.3x\ell = 9x$. A consumer's utility for such a contract x at a price p would depend on his type as follows:

$$U(p,x,L) = -[0.2e^{p+(1-x)30/10} + (1-0.2)e^{p/10}], \quad U(p,x,H) = -[0.6e^{p+(1-x)30/10} + (1-0.6)e^{p/10}].$$

Then the prices for a contract y to make a consumer of each type indifferent between y and the contract x at price p can be computed by the formulas:

$$\varphi(p,x,L,y) = p + 10 \text{LN}(0.2e^{(1-x)30/10} + 1-0.2) - 10 \text{LN}(0.2e^{(1-y)30/10} + 1-0.2),$$

$$\varphi(p,x,H,y) = p + 10 \text{LN}(0.6e^{(1-x)30/10} + 1-0.6) - 10 \text{LN}(0.6e^{(1-y)30/10} + 1-0.6).$$

In Figure 1 below, the two green curves are indifference curves for the low-risk types, and the two red curves (with higher slopes) are indifference curves for the high-risk types. Of the three dotted lines in Figure 1, the lowest (green) line is where $p=C(x,L)$, the highest (red) line is where $p=C(x,H)$, and the middle (blue) line is the pooling-cost line where $p=0.75C(x,L)+0.25C(x,H)$.

The two triangles in Figure 1 show the separating contracts that are sold in the competitive equilibrium. The red triangle is at the full-insurance contract $x_H=1$ that high-risk types can buy at the price $p_H=C(1,H)=18$. Then the green triangle is at the contract x_L which satisfies the equation $\varphi(p_H,x_H,H,x_L) = C(x_L,L)$, and so this contract $x_L=0.322$ can be sold to the low types at the price $p_L=C(x_L,L)=1.93$ without diverting the high types from buying $x_H=1$ at the price p_H .

The blue square in Figure 1 denotes the contract x^* which, when priced at the pooling cost $p^*=0.75C(x^*,L)+0.25C(x^*,H)$, would be most preferred by the low-risk types; here $x^*=0.820$ and $p^*=7.38$. This contract x^* at this pooling price p^* would be strictly better for both types of consumers than what they get in the separating equilibrium, and so this (x^*,p^*) pooling plan is Pareto-superior to the competitive equilibrium. (Notice that the firms are just breaking even in any case.) But to have an equilibrium where all consumers bought this contract x^* at this price p^* , the equilibrium prices for any smaller contract $x < x^*$ would have to be priced along the thin green indifference curve for the low types that goes through (x^*,p^*) , and these prices would be greater than the low-types' cost $C(x,L)$ when $0.458 \leq x < x^*$. Thus, we cannot have a pooling equilibrium at (x^*,p^*) as long as firms have the option to offer contracts x between 0.458 and x^* (between the large blue square and the small green square) at competitive prices that would attract low types away from the pooling equilibrium.

However, if government regulations prevented firms from offering any contract x in the interval where $0.457 < x < x^*$, then there could be a Pareto-superior competitive equilibrium where everyone buys the contract x^* at the price p^* . The virtual price for an untraded contract $x \leq 0.457$ would be $p = \varphi(p^*, x^*, L, x)$ on the thin green indifference curve for low types, and the virtual price for the untraded contracts $x > x^*$ would be $p = \varphi(p^*, x^*, H, x)$ on the thin red indifference curve for high types. In each case, these virtual prices would be the highest price that could just attract one type, but firms could not expect to profit by selling to the attracted type at these prices.




Figure 1.

Markets with adverse selection on the sellers' side (e.g. Spence signaling in labor markets)

We can also model markets where competitive firms buy labor or other inputs from individuals with private information about themselves that might be relevant for their potential productivity in any labor-supply contract. Again in this case we can let Y denote the set of possible labor contracts, and let T denote the set of possible types of individual workers. A given utility function $U: \mathbb{R} \times Y \times T \rightarrow \mathbb{R}$ specifies the utility $U(p, x, t)$ that any individual of type t would get from selling his labor under the terms of contract x at wage p . In this case, of course, the worker's utility $U(p, x, t)$ would be a strictly increasing function of the wage-price p . A given productivity function $V: Y \times T \rightarrow \mathbb{R}$ specifies the expected output value $V(x, t)$ for a firm from a labor contract x with a worker of type t . A given probability distribution μ in $\Delta(T)$ specifies the fraction $\mu(t) > 0$ of workers who have each type t in the general population. Each worker must enter into exactly one labor contract. (We could include the no-trade option as an $x=0$ contract with productivity 0 and utility 0 for all types.)

Our simple concept of competitive equilibrium can be directly extended to such models. Then a competitive equilibrium would be any pair (q, γ) such that $q = (q(x))_{x \in Y}$ is a price vector in \mathbb{R}^Y , $\gamma = (\gamma(x|t))_{x \in Y, t \in T}$ is an allocation vector in $\Delta(Y)^T$, and the following conditions are satisfied:

$$[0'] \quad \sum_{y \in Y} \gamma(y|t) = 1 \text{ and } \gamma(x|t) \geq 0, \forall x \in Y, \forall t \in T;$$

$$[1'] \quad \sum_{x \in Y} \gamma(x|t) U(q(x), x, t) = \max_{x \in Y} U(q(x), x, t), \forall t \in T;$$

$$[2'] \quad \sum_{t \in T} \mu(t) \gamma(x|t) (V(x, t) - q(x)) = 0, \forall x \in Y; \text{ and}$$

$$[3'] \quad \forall y \in Y, \exists t \in T \text{ such that } U(q(y), y, t) = \max_{x \in Y} U(q(x), x, t) \text{ and } q(y) \geq V(y, t).$$

Thus, for markets with adverse selection on the sellers' side, we can similarly require that [2'] competitive firms should expect to break even on any contract that attracts a positive fraction of the informed sellers, and that, [3'] even for an untraded contract, there should be at least one seller-type which would be willing to choose this contract but would not be profitable for the competitive firms at the given price.

Homework exercise on markets with adverse selection:

Consider a population of risk-averse consumers who face uncertain but potentially insurable losses that will be independently drawn from normal distributions. Each consumer has constant risk tolerance, with the risk-tolerance index $R=10$, and each consumer's loss will be drawn from a normal distribution with standard deviation $\sigma=20$. But the expected value of this loss will depend on the consumer's risk type, which may be high or low. A consumer of the low-risk type has an expected loss of $m_L=40$, but a consumer of the high-risk type has an expected loss of $m_H=60$. In the overall population, $\mu(L)=80\%$ of consumers are the low-risk type and $\mu(H)=20\%$ of consumers are the high risk type. Each consumer knows his own risk type, but insurance companies have no way to directly observe anyone risk type (other than by the different choices that they might make in the market).

Consider linear insurance contracts which are parameterized by the fraction x of the consumer's loss that the insurance company will cover, where $0 \leq x \leq 1$. That is, an x -contract specifies that the insurance company will pay $x\tilde{S}$ to the consumer when the consumer's risky loss turns out to be \tilde{S} . (With normally distributed losses, there is a small probability that a consumer's "loss" \tilde{S} might be negative, in which case $|\tilde{S}|$ would actually denote a risky income for the consumer, of which the consumer would be contractually obligated to pay the fraction $x|\tilde{S}|$ to the insurance company.)

(a) Find the coverage fraction \bar{y} such that a high-risk consumer would be indifferent between [1] buying full insurance ($x=1$) at a price such that risk-neutral insurers would expect to just break even in selling to high-risk consumers ($C(1,H)$) and [2] buying a \bar{y} -contract at a price such that risk-neutral insurers would expect to just break even in selling to low-risk consumers ($C(\bar{y},L)$).

(b) Construct a competitive equilibrium of this insurance market, when competitive insurance companies can sell linear insurance contracts for any fraction x between 0 and 1. You should specify equilibrium prices of x -contracts, for all x in the interval $[0,1]$, and you should identify the contracts that will be purchased by each type of consumer. Then verify that this equilibrium satisfies the equilibrium conditions that we introduced in class (the simplified version of Azevedo & Gottlieb's equilibrium concept). Also, compute the certainty-equivalent value $\bar{U}(q,t)$ for each type of consumer in this equilibrium.

(c) Now consider what contract the low-risk consumers would want to buy if all contracts could be priced so that risk-neutral insurers would expect to just break even in selling to the overall population (80% low-risk types and 20% high-risk types). Find the coverage fraction y^* that would maximize the certainty equivalent for a low-risk consumer in such an ideal world where contracts could be priced according to their expected cost with pooling of all consumers. Show that both the low-risk types and the high-risk types would strictly prefer to buy the contract y^* at this pooling price ($\sum_{t \in \{L,H\}} \mu(t)C(y^*,t)$) over the contracts that they buy with the equilibrium prices that you found in part (b).

(d) Find the coverage fraction \hat{y} such that a low-risk consumer would be indifferent between [1] buying a \hat{y} contract at the price such that risk-neutral insurers would expect to just break even in selling \hat{y} -contracts to low-risk consumers ($C(\hat{y},L)$) and [2] buying the y^* contract (from part (c)) at the pooling price such that risk-neutral insurers would expect to just break even in selling y^* -contracts to the whole population.

(e) Now suppose that the government banned the sale of insurance contracts with any coverage fraction x such that $\hat{y} < x < y^*$. Show that there is now a competitive equilibrium with pooling at y^* . Verify all our conditions for an equilibrium, and show that this equilibrium makes all consumers strictly better off than the competitive equilibrium that you found in part (b).

These notes are available at:

<http://home.uchicago.edu/~rmyerson/teaching/eqmadsel.pdf>