Signalling and Default: Rothschild–Stiglitz Reconsidered

Pradeep Dubey and John Geanakoplos*

Abstract

In our previous paper we built a general equilibrium model of default and punishment in which equilibrium always exists and endogenously determines asset promises, penalties, and sales constraints. In this paper we interpret the endogenous sales constraints as equilibrium signals. By specializing the default penalties and imposing an exclusivity constraint on asset sales, we obtain a perfectly competitive version of the Rothschild–Stiglitz model of insurance. In our model their separating equilibrium always exists even when they say it doesn't.

Keywords: default, incomplete markets, adverse selection, moral hazard, equilibrium refinement, signalling, endogenous assets

*The first version of this paper was contained in a manuscript circulated in 1994, and then printed as a Cowles Foundation Discussion Paper in 2000. The current version is a refinement of the second half of that paper.

1 Introduction

In our previous paper (Dubey, Geanakoplos, Shubik, 2001) we built a model that explicitly allows for default, but is broad enough to incorporate conventional general equilibrium theory as a special case. We call the model $GE(R, \lambda, Q)$ because each asset j is defined by its promise R_j, the penalty λ_j for default on the promise, and the quantity restriction Q_j attendant on those who sell it.

In the model, a seller h of φ_j^h units of asset j has the *option* of delivering any $D_j^h \leq \varphi_j^h R_j$, and incurring a penalty $\lambda_j^h[\varphi_j^h R_j - D_j^h]^+$ on the shortfall. As a result of the option, different sellers may pay off differently on the same asset. We maintain the hypothesis of perfect competition by supposing that buyers do not trade with individual sellers, but with the market. The buyers of asset j receive a pro rata share of all its different sellers' deliveries, just as an investor does today in the securitized mortgage market. Thus from the buyer's perspective, asset j is a share in a pool of deliveries. Each unit purchased of the asset delivers

$$\delta_j \equiv \frac{\sum_h D_j^h}{\sum_h \varphi_j^h}.$$

The buyer of asset j does not know the identity of its sellers, nor the quantities of their sales. But he does know that none of them could have sold more than the limit Q_j.

The pooling leads to adverse selection, since a buyer must worry that sellers with a proclivity for default (on account of low default penalties or low endowments) will tend to sell more of the asset, worsening the anticipated rate of delivery δ_j. Moral hazard enters the picture because a seller can choose to default, and because an agent who sells more promises will be less able to fully deliver on any one of them.

Signalling, by publicly committing oneself to a small quantity of sales, therefore has an important role to play, because it assures the buyer of a more reliable delivery. Two assets i and j with identical promises $R_i = R_j$ and penalties $\lambda_i = \lambda_j$, may sell for different prices $\pi_i > \pi_j$ if $Q_i < Q_j$. Conforming mortgages, which are presently limited to $275,000, sell for a higher price (per dollar promised) than jumbo mortgages, which are not so constrained in size.

Without any default, there is no option, and our model reduces to the standard general equilibrium model with incomplete markets (GEI). Pooling different deliveries is what distinguishes our model from GEI, and enables it to include phenomena like adverse selection and moral hazard that are missing from GEI.

Pooling, however, does not compromise the existence of equilibrium. We showed in our previous paper that for any exogenously fixed set $\mathcal{A}$ of tradeable assets,

$$\mathcal{A} = \{(R_j, \lambda_j, Q_j) : (R_j, \lambda_j, Q_j) \text{ is tradeable}\},$$

equilibrium $E(\mathcal{A})$ always exists. The levels of trade, the rates of default, and the prices of all assets in $\mathcal{A}$ emerge endogenously as part of the equilibrium $E(\mathcal{A})$. In contrast to standard general equilibrium models, in which prices are the only equilibrating variables, here prices and anticipated delivery rates are needed to clear markets.

A key feature of our equilibrium is a condition on expected deliveries of untraded assets that is similar to the trembling hand refinements used in game theory. Notice that δ_j is not defined when total sales $\sum_h \varphi_j^h = 0$. The condition says that δ_j should be derived as the limit of delivery rates $\delta_j(\varepsilon)$ taken over a sequence of small perturbations $\varepsilon \to 0$, at which there is positive trade of asset j.

Equilibrium also gives rise to the subset $\mathcal{A}^* \equiv \mathcal{A}^*(E(\mathcal{A})) \subset \mathcal{A}$ of actively traded assets:

$$\mathcal{A}^* = \{(R_j, \lambda_j, Q_j) \in \mathcal{A} : (R_j, \lambda_j, Q_j) \text{ is positvely traded in } E(\mathcal{A})\}.$$

Using the refinement, we argued in our previous paper that equilibrium endogenously determines the traded assets; with default, $\mathcal{A}^*$ tends to be much smaller than $\mathcal{A}$.

This is in sharp contrast to GEI, which is a special case of our model, with $\lambda = Q = \infty$. In GEI, for generic utilities and endowments, the span of the actively traded promises $\mathcal{A}^*$ equals the span of the available promises $\mathcal{A}$. Thus effectively $\mathcal{A}^* = \mathcal{A}$, and GEI is unable to explain the endogeneity of traded assets.[1]

In this paper our focus is on the endogeneity of the quantity signals Q_j. To this end, we fix the penalties $\lambda_j = \bar{\lambda}$ and promises $R_j = \bar{R}$ and see which sales restrictions Q_j emerge in $\mathcal{A}^*$. If the quantity constraints Q_j on asset sales are not binding in equilibrium, then there is in effect no signalling. Otherwise the Q_j become signals which play a crucial role in the equilibrium. The phenomenon of signalling can thus be treated in perfect competition, moreover without jeopardizing the existence of equilibrium.

By suitable choices of default penalties we subsume insurance contracts in our framework. Take

$$\lambda_{sj}^h = \begin{cases} \infty & \text{if} \quad s \in \bar{S}^h \subset S \\ 0 & \text{if} \quad s \notin \bar{S}^h \end{cases} .$$

Agents are forgiven completely in some states (perhaps when their endowments are zero) and compelled to repay otherwise. This enables us to capture insurance in terms of trading assets in our model. Consider an asset which promises one dollar in every state, but whose expected delivery rate is $\delta < 1$ on account of default. An agent h, who buys and sells one unit of this asset, will fully deliver in his good states (since $\lambda_s^h = \infty$ for $s \in \bar{S}^h$) and fully default in his bad states (since $\lambda_s^h = 0$ for $s \notin \bar{S}^h$). On net, h then obtains δ in his bad states by giving up $1 - \delta$ in his good states, which is tantamount to taking out insurance. In particular, the models of Akerlof (1972) and Rothschild and Stiglitz (1976) can be embedded as special cases in our model.

Rothschild and Stiglitz seem to have had in mind oligopolistic insurance companies, designing contracts for a continuum of private agents. These companies intermediated trade between the agents, setting prices to attract customers, and standing

[1] Only when transactions costs for assets are introduced into the GEI model can $\mathcal{A}^*$ substantively differ from $\mathcal{A}$. See for example the work of Allen–Gale (1988) and Pesendorfer (1995).

ready to insure all who accepted their offers. They were *assumed* to have risk-neutral preferences and to maximize expected profits.

We have recast this story in a perfectly competitive setting, retaining only the continuum of agents. We do not have insurance companies — we have markets. Diverse groups of agents trade promises through these massive, anonymous markets. Since the assets bought are pools of promises, and those sold permit idiosyncratic deliveries and default, the net effect is that agents insure each other through the markets. *Every* agent is a price taker. Yet the model is subtle enough to unambiguously determine which insurance contracts will emerge in $\mathcal{A}^*$. In our model, the market forces of perfect competition take over the role of designing contracts. Furthermore we *derive* the conclusion that insurance policies yield zero expected profit, solely from the distribution of accident risk in the economy, without postulating any risk-neutral agents.

Most remarkable is that our existence theorem for equilibrium with default also guarantees the existence of insurance equilibrium, in spite of the adverse selection. Rothschild and Stiglitz showed that there were robust economies in which their equilibrium does not exist. By replacing a hybrid competitive/oligopolistic model with a simpler perfectly competitive model, we are able to retain the subtlety of the separating equilibrium, and at the same time to restore the universal existence of equilibrium.

In Section 2 we recall our basic model of default and we discuss the equilibrium refinement. Section 3 restates two existence theorems proved in our previous paper. In Section 4 we show how to reinterpret the variables in the model so as to include insurance as a special case. In Section 5 we illustrate the emergence of adverse selection in our model, in contrast to GEI. When there is only one available asset (with no limitation on its sale), agents of different types will sell it. Unreliable agents, for whom it is less costly to default, will sell disproportionately large quantities, adversely affecting delivery rates on the pool.

In Section 6 we consider signalling in the context of insurance, i.e., we allow for many assets with identical promises and penalties, but different sales limits. We further specialize the model, confining ourselves to infinite/zero penalties and to two types of agents (reliable and unreliable), to bring it into the Rothschild–Stiglitz framework. Like them, we presume that asset sales are exclusive, that is that each agent can sell at most one asset (i.e., obtain at most one insurance contract). As shown by Theorem 2 in Section 3, this exclusivity constraint does not endanger the existence of equilibrium. We show that their "separating" contracts always form a $GE(R, \lambda, Q)$ equilibrium, even when they say there is no equilibrium. Moreover, our equilibrium refinement is stringent enough to guarantee that this is the unique equilibrium.

The crucial difference between the Rothschild–Stiglitz definition of equilibrium and ours can be understood in terms of the assumption each makes about the reliability of untraded contracts $j \in \mathcal{A} \backslash \mathcal{A}^*$. We argue in Section 2 and Section 6 that our assumption is natural when there are many buyers and sellers in perfect competition, and corresponds to cautious expectations (determined by the "trembling hand" refinement to equilibrium that we give). By contrast, the expectations attributed to agents by Rothschild and Stiglitz are not compatible (to our way of thinking) with perfect competition.

In our sequel paper (Dubey–Geanakoplos, 2001) we explore what happens in the more realistic setting where agents can sell more than one asset, albeit under restrictions. When the assets are ranked by seniority, we find that a new kind of primary–secondary insurance equilibrium emerges, which appears to conform better with real practice. Both the reliable and unreliable agents take out the same primary insurance policy, at intermediate rates; in addition, the unreliable take out more secondary insurance at an unfavorable rate.

2 Default in Equilibrium: The $GE(R, \lambda, Q)$ Model

2.1 The Economy

As in the canonical model of general equilibrium with incomplete markets (GEI), we consider a two-period economy, where agents know the present but face an uncertain future. In period 0 (the present) there is just one state of nature (called state 0), in which H agents trade in L commodities and J assets. Then chance moves and selects one of S states which occur in period 1 (the future). Commodity trades take place again, and assets pay off. The difference from GEI is that in our $GE(R, \lambda, Q)$ model, assets pay off in accordance with what agents opt to deliver. Our notation can be formalized as follows:

$\ell \in L = \{1, ..., L\}$ = set of commodities
$s \in S = \{1, ..., S\}$ = set of states in period 1
$S^* = \{0\} \cup S$ = set of all states
$h \in H = \{1, ..., H\}$ = set of agents
$e^h \in \mathbb{R}_+^{S^* \times L}$ = initial endowment of agent h
$j \in J = \{1, ..., J\}$ = set of assets
$R_j \in \mathbb{R}_+^{S \times L}$ = promises per unit of asset j of each commodity $\ell \in L$ in each state $s \in S$
$u^h : \mathbb{R}_+^{S^* \times L} \to \mathbb{R}$ = utility function of agent h
$\lambda_{sj}^h \in \overline{\mathbb{R}}_+ \equiv \mathbb{R}_+ \cup \{\infty\}$ = real default penalty on agent h for asset j in state s
$Q_j^h \in \mathbb{R}_+$ = bound on sale of asset j by agent h

We assume that no agent has the null endowment, and that all named commodities are present in the aggregate, i.e.,

$$e_s^h = (e_{s1}^h, ..., e_{sL}^h) \neq 0$$

for all $h \in H$ and $s \in S^*$, and

$$e_{s\ell} = \sum_{h \in H} e_{s\ell}^h > 0$$

for all $s\ell \in S^* \times L$. Also each u^h is continuous, concave and strictly increasing in each of its $S^* \times L$ variables. Having assumed strict monotonicity and concavity, there is no further loss of generality in assuming that $u^h(x) \to \infty$ whenever $||x||_\infty \to \infty$.[2]

We can visualize the state space as a simple tree:

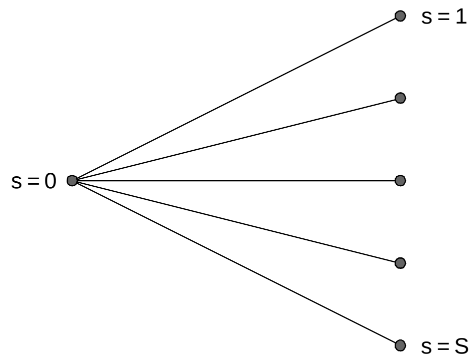

Figure 1

Agents h have heterogeneous, state-dependent endowments $e^h_s \in \mathbb{R}^L_+$ and disutilities of default λ^h_{sj}.

Adverse selection enters the picture because agents have different endowments out of which to keep their promises, and also different disutilities of default.

Promises must be of a limited kind $j \in J$ fixed a priori. A promise $j \in J$ specifies bundles of goods (or services) to be delivered in each state:

$$\text{Promise } R_j = \left(\begin{array}{c} \\ \\ \\ \end{array} \right) \begin{array}{l} \} - \text{state 1 goods} \\ \} - \text{state 2 goods} \\ \\ \} - \text{state } S \text{ goods.} \end{array}$$

Agents h make promises by selling various quantities φ^h_j of each asset j. An agent's ability to keep a promise depends on how many promises he sells, both of the same kind j, and of other kinds $j' \neq j$. Moral hazard enters the picture, since a buyer of an asset (i.e., lender) does not know which other promises the seller (i.e., borrower) has made, and because borrowers have the option to default.

Each kind of asset prescribes a limit on its sale, $\varphi^h_j \leq Q^h_j$. Limits on sales of promises are necessary to any realistic model of credit.[3] If $Q^h_j = 0$, then agent h is essentially forbidden from selling asset j. If the limits Q^h_j are very large, they may be entirely irrelevant. But if they are small, then they may be used as a signal that the sellers are not making many promises, and hence that the promises are reliable.

[2] Let $\Box = \{x \in \mathbb{R}^{S^* L}_+ : ||x||_\infty \leq 2||\sum_h e^h||_\infty\}$. Let $\mathcal{L}$ be the set of affine functions $L : \mathbb{R}^{S^* L}_+ \to \mathbb{R}$ such that $L(x) \geq u^h(x)$ for all $x \in \Box$. Define $\tilde{u}^h(x) \equiv \inf_{L \in \mathcal{L}} L(x)$. Then equilibrium with u^h and $\tilde{u}^h$ coincide, and $\tilde{u}^h$ has the desired properties.

[3] Evidence abounds that finite bounds are always imposed in the extension of credit. Even the best "name" among borrowers has a limited credit line.

An economy is defined as a vector

$$\mathcal{E} = \left((u^h, e^h)_{h\in H},\ \left(R_j, ((\lambda^h_{sj})_{s\in S}, Q^h_j)_{h\in H} \right)_{j\in J} \right).$$

Note again that an asset consists of promises, penalties for default, and limits on sales.

2.2 Equilibrium

To define a $GE(R,\lambda,Q)$ equilibrium, first consider the "macrovariables" p, π, K that each agent takes as fixed. Here $p \in \mathbb{R}^{S^*\times L}_{++}$ is the vector of commodity prices; $\pi \in \mathbb{R}^J_+$ is the vector of asset prices; and K is an $S\times J$ matrix with entries K_{sj} between 0 and 1, representing the fraction expected to be delivered of payments promised by asset j in state s.

2.2.1 Household Budget and Payoff

The *budget set* $B^h(p,\pi,K)$ of agent h is given by:

$$\begin{aligned} B^h(p,\pi,K) \;=\; \Big\{ & (x,\theta,\varphi,D) \in \mathbb{R}^{S^*\times L}_+ \times \mathbb{R}^J_+ \times \mathbb{R}^J_+ \times \mathbb{R}^{J\times S\times L}_+ : \\ & p_0\cdot(x_0 - e^h_0) + \pi\cdot(\theta-\varphi) \le 0;\ \varphi_j \le Q^h_j \text{ for } j\in J;\ \text{and, } \forall s\in S, \\ & p_s\cdot(x_s - e^h_s) + \sum_{j\in J} p_s\cdot D_{sj} \le \sum_{j\in J} \theta_j K_{sj} p_s\cdot R_{sj} \Big\} \end{aligned}$$

Here $x \in \mathbb{R}^{S^*\times L}_+$ is the final consumption of commodities, $\theta\in\mathbb{R}^J_+$ (respectively, $\varphi\in\mathbb{R}^J_+$) gives the purchases (respectively, sales) of the J assets, and $D_{sj}\in\mathbb{R}^L_+$ is the vector of goods delivered by agent h on asset j in state s.

The budget set allows agent h to deliver whatever he pleases. On the other hand, the agent expects to receive a fraction K_{sj} of the promises made to him on asset j in state s. The first constraint says that agent h cannot spend more on purchases of commodities x_0 and assets θ than the revenue he receives from the sale of commodities e^h_0 and assets φ. Moreover he can never sell more than Q^h_j of any asset j. The second constraint applies separately in each state $s\in S$. It says that agent h cannot spend more on the purchase of commodities x_s and asset deliveries $\sum_j D_{sj}$ in state s than the revenue he gets in state s from commodity sales e^h_s and asset receipts $\sum_j \theta_j K_{sj} p_s R_{sj}$.

The only reason that agents deliver anything on their promises is that they feel a disutility λ^h_{sj} from defaulting. The payoff of (x,θ,φ,D) given prices p, to agent h is

$$w^h(x,\theta,\varphi,D,p) = u^h(x) - \sum_{j\in J}\sum_{s\in S} \frac{\lambda^h_{sj}[\varphi_j p_s\cdot R_{sj} - p_s\cdot D_{sj}]^+}{p_s\cdot v_s}.$$

(Here $v_s \in \mathbb{R}^L_+ \backslash \{0\}$ represents a fixed basket of goods, using which default can be measured in real terms.) Note that $[\varphi_j p_s \cdot R_{sj} - p_s \cdot D_{sj}]^+ \equiv \max\{0, \varphi_j p_s \cdot R_{sj} - p_s \cdot D_{sj}\}$ is exactly the money value of the default of h on his promise to deliver on asset j in state s.

Notice that the budget set is convex, and the payoff function w^h is concave, in the household choice variables (x, θ, φ, D). Had we expressed these choices with other (apparently natural) variables, such as $\delta^h_{sj} \equiv$ delivery per unit promised, the budget set would no longer be convex, nor would w^h be concave.

It is worth noting a *scaling property* of the budget set (which is immediate from its definition and the fact that $e^h_s \neq 0$ and $p_s \gg 0$ for all $s \in S^*$): $(x, \theta, \varphi, D) \in B^h(p, \pi, K)$ and $0 < \alpha < 1 \Rightarrow (\alpha x, \alpha\theta, \alpha\varphi, \alpha D) \in B^h(p', \pi', K')$ for all (p', π', K') sufficiently close to (p, π, K). This property will often be useful to us.[4]

2.2.2 Market Clearing

We are now in a position to define a $GE(R, \lambda, Q)$ equilibrium. It is a list $\langle p, \pi, K, (x^h, \theta^h, \varphi^h, D^h)_{h \in H} \rangle$ such that (1) to (4) below hold.

(1) For $h \in H$, $(x^h, \theta^h, \varphi^h, D^h) \in \arg\max w^h(x, \theta, \varphi, D, p)$ over $B^h(p, \pi, K)$

(2) $\sum_{h \in H}(x^h - e^h) = 0$

(3) $\sum_{h \in H}(\theta^h - \varphi^h) = 0$

(4) $K_{sj} = \begin{cases} \sum_{h \in H} p_s \cdot D^h_{sj} / \sum_{h \in H} p_s \cdot R_{sj}\varphi^h_j, & \text{if } \sum_{h \in H} p_s \cdot R_{sj}\varphi^h_j > 0 \\ \text{arbitrary,} & \text{if } \sum_{h \in H} p_s \cdot R_{sj}\varphi^h_j = 0 \end{cases}$

Condition (1) says that all agents optimize; (2) and (3) require commodity and asset markets to clear. Condition (4), together with the definition of the budget set, says that each potential lender (i.e., buyer) of an asset is correct in his expectation about the fraction of promises that do in fact get delivered. Moreover, his expectation $K^h_{sj} = K_{sj}$ of the rate of delivery does not depend on anything he does himself; in particular, it does not depend on the amount θ^h_j he loans (i.e., purchases) of the asset. Every lender gets the same rate of delivery.

Since heterogeneous borrowers may be selling the same asset, the realized rate of delivery K_{sj} is an average of the rates of delivery of each of the borrowers, weighted by the quantity of their sales. It might well happen that those borrowers with the highest rates of default are selling most of the asset, and this is the adverse selection and moral hazard that rational lenders must forecast.

We believe that our definition of $GE(R, \lambda, Q)$ equilibrium embodies the spirit of perfect, anonymous competition, and represents a significant fraction of the mass asset markets of a modern enterprise economy.

In the next sections we investigate the properties of equilibrium.

[4] An alternative scaling property, also satisfied by the budget set, is obtained if we replace $(\alpha x, \alpha\theta, \alpha\varphi, \alpha D)$ with $(\alpha x, \alpha\theta, \varphi, \alpha D)$. Our entire analysis remains intact with this version of scaling.

2.3 An Equilibrium Refinement

When assets are traded, expected deliveries K_{sj} must be equal to actual deliveries. Expectations cannot therefore be unduly pessimistic. But for assets that are not traded, our model so far makes no assumption about expectations of delivery (see (4)).

We believe that unreasonable pessimism prevents many real world markets from opening, and provides an important role for government intervention. But it is interesting to study equilibrium in which expectations are always reasonably optimistic. It is of central importance for us to understand which markets are open and which are not, and we do not want our answer to depend on the agents' whimsical pessimism. Without further conditions there are always trivial equilibria in which all $\pi_j = 0$ and $K_{sj} = 0$. Since $\pi_j = 0$, nobody wants to sell, and since $K_{sj} = 0$, nobody wants to buy. To avoid this we add a condition (5) to the definition of equilibrium. This requires that if a small change in the macro parameters (p, π) could induce some agents to start selling some of an asset j, where none was being sold before, then buyers should expect at least the rate of delivery they would get had the world indeed been so perturbed. (If there are many ways of perturbing (p, π) to induce sales, then we allow the buyers to focus their attention on one of these perturbations.) If prices π_j are so low that no small perturbation will induce any agents to sell asset j, then buyers are required to expect full delivery, $K_{sj} = 1$. One can (but need not) interpret these expectations as if the government guaranteed delivery on the first infinitesimal promises. In our previous paper, we showed that condition (5) can always be realized by adding an extra agent to the economy who sells ε of every promise and always fully delivers on his promises, and then letting $\varepsilon \to 0$.

Let $\| \|_\infty$ denote the supremum norm, and let $E \equiv \langle p, \pi, K, (x^h, \theta^h, \varphi^h, D^h)_{h\in H}\rangle$ be a candidate equilibrium which satisfies conditions (1) to (4). For $s \in S$, let $J(s) = \{j \in J : \sum_{h\in H} p_s \cdot R_{sj}\varphi_j^h = 0\}$. Thus $J(s)$ is the set of assets in state s for which K_{sj} is not determined by market activity in E. We are ready to state

(5) For any $\varepsilon > 0$, there exists $E(\varepsilon) \equiv \langle p(\varepsilon), \pi(\varepsilon), K(\varepsilon), (x^h(\varepsilon), \theta^h(\varepsilon), \varphi^h(\varepsilon), D^h(\varepsilon))_{h\in H}\rangle$ such that

- (i) $(x^h(\varepsilon), \theta^h(\varepsilon), \varphi^h(\varepsilon), D^h(\varepsilon)) \in \arg\max w^h(x, \theta, \varphi, D, p(\varepsilon))$ over $B^h(p(\varepsilon), \pi(\varepsilon), K(\varepsilon))$
- (ii) $\|E - E(\varepsilon)\|_\infty < \varepsilon$
- (iii) $$K_{sj}(\varepsilon) \geq \begin{cases} \sum_{h\in H} p_s(\varepsilon)\cdot D_{sj}^h(\varepsilon) / \sum_{h\in H} p_s(\varepsilon)\cdot R_{sj}\varphi_j^h(\varepsilon) & \text{if } \sum_{h\in H} p_s(\varepsilon)\cdot R_{sj}\varphi_j^h(\varepsilon) > 0 \\ 1 & \text{if } \sum_{h\in H} p_s(\varepsilon)\cdot R_{sj}\varphi_j^h(\varepsilon) = 0 \end{cases}$$

  for all $s \in S$ and $j \in J(s)$.

Conditions (i), (ii), and (iii) say that if asset j is untraded and $K_{sj} < 1$, then there must be arbitrarily small perturbations of the macro variables which induce agents to sell j, and to deliver (in aggregate) at a rate at most K_{sj} (that is, to default at rate at least $1 - K_{sj}$). If $K_{sj} = 1$, the refinement is automatically satisfied.

Our condition (5) will enable us to ascertain when an asset, though priced, is not traded in equilibrium.

Consider the following heuristic example, illustrated in Figure 2 below. Suppose an equilibrium in an economy with assets $j = 1, ..., J$ is given. A new asset $J+1$, promising 1 in every state, is added to the economy. Suppose at prices $\pi_{J+1} \leq 10$, no agent would sell it, while at prices $10 < \pi_{J+1} \leq 20$ only unreliable agents (with low default penalties $\lambda^h_{s,J+1}$) would sell it, and at prices $\pi_{J+1} > 20$ reliable and unreliable agents would sell it. Without the equilibrium refinement, we could always include the new asset in the old equilibrium by assigning it a price $\pi_{J+1} = 0$ with no trade, and with $K_{sJ+1} = 0 \ \forall s \in S$. But our equilibrium refinement requires that if $\pi_{J+1} < 10$, then $K_{sJ+1} = 1$ for all $s \in S$, since no perturbation would induce sales. The equilibrium refinement thus rules out equilibria with $\pi_{J+1} < 10$ unless demand is zero even with expectations of full delivery. For concreteness, let us suppose demand is zero unless expected delivery per dollar invested is at least 0.034, after which demand becomes positive. Clearly there is no equilibrium with $0 \leq \pi_{J+1} < 10$, since expected delivery per dollar invested $\frac{1}{S}\sum_s K_{sj} 1/\pi_{J+1} = (1)(1/\pi_{J+1}) \geq (1)(1/10) = 0.100 > 0.034$.

If there is any equilibrium in this example in which asset $J+1$ is not traded, then there must be such an equilibrium at which $\pi_{J+1} = 10$. Since there are no sales, when $\pi_{J+1} = 10$, the refinement allows for $K_{sJ+1} < 1$, provided that there would be sales at $\pi_{J+1} = 10 + \varepsilon$ and that the delivery rate on those sales is approximately K_{sJ+1}, or lower. By hypothesis, at $\pi_{J+1} = 10 + \varepsilon$, only unreliable agents would be selling. Suppose unreliable agents always deliver 1/3 of what they promise. Then we might have $K_{sJ+1} = 1/3 \ \forall s \in S$. At these low levels of delivery, and at a price of 10, there would indeed be no buyers (as well as no sellers), since expected delivery per dollar invested in $(1/3)(1/10) = 0.033 < 0.034$. Since the K_{sJ+1} are obtained from the perturbation, we would regard the expectations as reasonable and call this a genuine equilibrium.

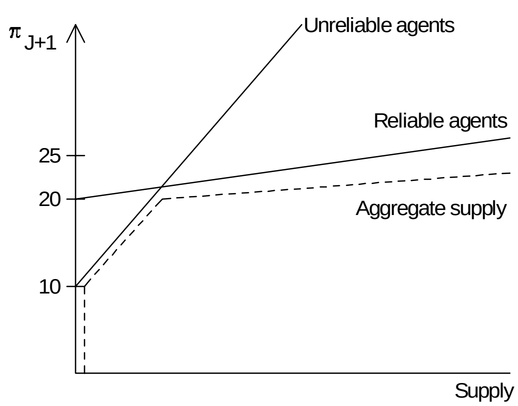

Figure 2

The Rothschild and Stiglitz logic would say that one must take into account the delivery rates of the reliable agents. Suppose, for concreteness, that they deliver $K'_{sJ+1} = 9/10$ no matter how much they promise. Further, suppose that if the price were raised to $\pi_{J+1} = 25$, reliable agents would want to sell in such large quantities that, leaving aside the question of market clearing, the fraction of deliveries out of all desired sales on asset $J+1$ would be $K'_{sJ+1} = 6/7$. (If at price $\pi_{J+1} = 25$, reliable agents sell 110 units to every 9 units unreliable agents sell, then $K'_{sJ+1} = 6/7 = (110/119)(9/10) + (9/119)(1/3)$.)

If buyers took 6/7 as the rate of delivery, then at the price $\pi_{J+1} = 25$ they would be willing to buy, even though they had refused to buy at $\pi_{J+1} = 10$, because their returns per dollar invested would be better, $6/7 \cdot 1/25 > (0.34)(1/10) > 1/3 \cdot 1/10$. If this were the case, then according to the logic of Rothschild–Stiglitz (1972), equilibrium would not exist at $\pi_{J+1} = 10$, since buyers would have an incentive to raise the price to $\pi_{J+1} = 25$.

Our definition of equilibrium allows for $\pi_{J+1} = 10$, and $K_{sJ+1} = 1/3$, and we believe it does so for good reasons. First, we suppose that buyers are aware of the composition of sales at the market prices, and perhaps of the composition of sales at prices a penny off from market prices. But agents lack the knowledge or computing power to infer what the composition would be at prices far from market prices.Second, we have in mind a competitive world with many small buyers. If a single buyer raised his offering price to 25, fully 15 points above the market price, he would be deluged with sellers. The people with the most to gain from selling to him would be those who already were willing to sell at 10, namely the unreliable agents. Why should he assume he would be equally likely to encounter each unit sold? We feel justified in assigning him the cautious expectations of $K_{sJ+1} = 1/3$ no matter what price he offered, given that the market price is $\pi_{J+1} = 10$.

In our model, agents do not unilaterally set prices; they are price takers and the market sets the price. We regard an asset or contract as setting out the obligations of the seller, including the penalties if he fails to deliver, and the quantity limitations on his other sales. The price of the contract is set by competition between sellers and buyers, that is, by the market. Agents need only think about one prevailing price for each contract. In our view, competitive equilibrium should be defined by a single price at which both supply and demand are equal (possibly both zero, as long as expectations at that price are set at rational levels).

In the Rothschild–Stiglitz view, the price is one of the terms of the contract. In this view, there is no such thing as a single contract; there are as many contracts as there are prices. Notice also that the Rothschild–Stiglitz view must regard market clearing as one of rationing. At most prices, the contract will not be traded, because *either* supply or demand is zero, and the other side of the market is rationed. This point of view has been admirably expressed by Gale.

The ability to unilaterally set prices, and to compute what demand would be at these different prices, are features of oligopolistic models. We rigorously maintain the hypothesis of perfect competition, which rules out both these features. Our equilibrium refinement nevertheless has bite, by ruling out many no-trade equilibria, even in our perfectly competitive framework. Consider the situation where the unreliable and reliable supply curves are reversed, so that at $\pi_{J+1} = 10$ it is the reliable agents who begin to sell. Then according to our definition of equilibrium, in order for $\pi_{J+1} = 10$ to be an equilibrium, it must be that no demand would be forthcoming even with K_{sJ+1} set at the reliable rates of delivery ($K_{sJ+1} = 9/10$). But $(9/10)(1/10) = 0.090 > 0.034$, so in this version of the example, there could be no equilibrium in which asset $J+1$ remains untraded. (Our existence theorem, stated in the next section then assures us that there must be some equilibrium in which asset $J+1$ is traded.)

2.4 A Continuum of Traders

We have mentioned several times that our model is meant to embody the ideal of perfect competition, in which each agent is so small that by himself he cannot influence anyone else. We can make such an interpretation of our model more concrete by replacing each agent h by a continuum of identical agents parameterized by t lying in the interval $(h-1, h]$: each agent $t \in (h-1, h]$ has identical characteristics: $(e^t, u^t, \lambda^t, Q^t) \equiv (e^h, u^h, \lambda^h, Q^h)$.

For any $(p, \pi, K) \in \mathbb{R}_{++}^{S^* \times L} \times \mathbb{R}_+^J \times [0,1]^{S \times J}$ we can define $B^t(p, \pi, K)$ exactly as before, replacing h by t throughout. Also $GE(R, \lambda, Q)$ can be defined as before, replacing $\sum_{h \in H}$ by $\int_I d\mu$, where $\mu \equiv$ Lebesgue measure, "$\forall h \in H$" by "almost all $t \in I$," and the notion of convergence of x, θ, φ, D (which are now integrable functions on I) in condition (5) by almost everywhere pointwise convergence.[5]

The $GE(R, \lambda, Q)$ of the finite agent economy, whose existence we shall prove in Theorem 1, corresponds to a $GE(R, \lambda, Q)$ of the continuum model with the added feature that $(x^t, \theta^t, \varphi^t, D^t) = (x^{t'}, \theta^{t'}, \varphi^{t'}, D^{t'})$ whenever t and t' are both in $(h-1, h]$, i.e., all agents of the same type behave symmetrically. We shall call such equilibria *type-symmetric* when viewed in the continuum setting.

But we shall shortly consider a variant of our model in which the convexity of budget sets fails to hold. Here the continuum model is necessary for establishing existence of $GE(R, \lambda, Q)$. Even if the economy is finite-type, its equilibria need no longer be type-symmetric, and the consideration of a continuum becomes unavoidable. See Theorem 2.

3 The Orderly Function of Markets with Default

In our previous paper we established that default is completely consistent with the orderly function of markets. To that end we proved that under fairly general conditions, equilibrium always exists in our model.

Theorem 1 *For any $\lambda \in \overline{\mathbb{R}}_+^{HSJ}$ and $Q \in \mathbb{R}_+^{RJ}$, a $GE(R, \lambda, Q)$ equilibrium satisfying* (1)–(5) *exists.*

The universal existence of equilibrium is somewhat surprising because of the historical tendency to associate default with disequilibrium (or more accurately, to make

[5] Without too much more trouble we could have allowed for an infinity of types. We have made the finite-type assumption only for ease of exposition.

full delivery part of the definition of equilibrium), as we have already remarked. Furthermore, endogeneity of the asset payoff structure is known to complicate the existence of equilibrium with incomplete markets. But we showed that no new existence problems arise from the endogeneity of the asset payoffs due to default.

So far our budget set is convex. But the simple quantity constraints we have already introduced do not allow us to formalize a wide enough variety of signals. They cannot handle cases when the sales constraints interact across assets j. In particular, they cannot handle exclusivity, under which an agent is prohibited from selling more than one asset. To take such constraints into account, we need to consider nonconvex budget sets $\underline{B}^h(p, \pi, K) \subset B^h(p, \pi, K)$.

We then require that the correspondences $\underline{B}^h$, $h \in H$, satisfy:

(0*) $(e^h, 0, 0, 0) \in \underline{B}^h(p, \pi, K) \subset B^h(p, \pi, K)$

(1*) $\underline{B}^h$ is upper semi-continuous

(2*) $\underline{B}^h$ has the scaling property: $(x, \theta, \varphi, D) \in \underline{B}^h(p, \pi, K)$ and $0 \leq \alpha < 1 \Rightarrow (\alpha x, \alpha\theta, \alpha\varphi, \alpha D) \in \underline{B}^h(\hat{p}, \hat{\pi}, \hat{K})$ for $(\hat{p}, \hat{\pi}, \hat{K})$ sufficiently close to (p, π, K)

The following theorem is a corollary of Theorem 7 in our previous paper.

Theorem 2 *Define equilibrium with budget sets* $\underline{B}^h(p, \pi, K)$ *for* $t \in (h-1, h]$ *and* $h \in H$*, satisfying* (0*), (1*), (2*) *above. Suppose the quantity constraints* Q_j^h *are all finite. Then, in the finite-type continuum model, equilibrium exists (though it may not be type symmetric).*

The universal existence of equilibrium with default is also surprising because the pioneering papers placing adverse selection in a model of competition, by Akerlof (1972) on the market for lemons, and Rothschild and Stiglitz (1976) on insurance markets, purportedly showed that adverse selection is quite commonly inconsistent with equilibrium.

Insurance contracts promise payments conditional on the state of nature, and so can be viewed as assets such as we describe in this paper. In particular, the Rothschild–Stiglitz model can be expressed as a special case of our general equilibrium model, as we show in Sections 4 and 5. The reason Rothschild and Stiglitz found robust regions with no equilibrium is that they defined equilibrium expectations differently, as we have explained. If buyers had the perfectly competitive expectations that we invoke, namely that each thinks he cannot improve his selection of sellers by unilaterally offering a higher price, then the Rothschild–Stiglitz model would always have an equilibrium, as we show in Section 5, even using their "exclusivity" hypothesis.

4 Default and Insurance

4.1 The Insurance Problem

As in Rothschild–Stiglitz, we consider a continuum of two types of agents: "reliable" (R) and "unreliable" (U). Each agent knows his own type, but not that of the others. Each agent has wealth (for simplicity, 1 dollar) in his "good" (no-accident) state, but nothing in his "bad" (accident) state for which he seeks insurance. Accidents occur independently across agents. The unreliable agents are more accident-prone than the reliable. Thus if $\text{prob}^t(G)$ denotes the probability of a good state for type t, we have $\text{prob}^R(G) > \text{prob}^U(G)$.

The utility for x units of money is $u(x)$, invariant of the state as well as agent-type. As is standard, we assume that u is strictly concave, monotonic and $u'(x) \to \infty$ as $x \to 0$. The consumption of (x_G, x_B) across the two states yields expected utility

$$\text{prob}^t(G)u(x_G) + \text{prob}^t(B)u(x_B)$$

to type $t = R, U$. For ease of presentation we take $\text{prob}^t(G)$ to be rational, and we further suppose that there is an equal population of each type.

To enable exact computation, we will focus on the numerical case $u \equiv \log$, $\text{prob}^U(G) = 1/3$, $\text{prob}^R(G) = 2/3$; but our analysis holds verbatim for the general scenario.

4.2 A Microeconomic Representation of Insurance

We recast the Rothschild–Stiglitz story into our framework, building a microfoundation for the insurance problem in the process. The key step is to represent probability distributions of accidents by states of the world which make explicit who has an accident there. This makes it clear that "identical" insurance policies for two agents of the same type do not pay off identically, since the agents will have accidents in different states, even if their probabilities are the same.

Within our framework of finite states and agents, we cannot maintain both the hypotheses that accidents occur independently, and that the same proportion of each type has an accident in every state. We drop the independence hypothesis, which actually plays no role in the theory anyway.

Since probabilities are rational, let $\text{prob}^R(B) = r/n$ and let $\text{prob}^U(B) = u/n$. To convert the insurance problem into our framework, take $\#S = n$, and suppose there are $\binom{n}{r}$ subtypes of reliable agents, each with population measure $\binom{n}{u}\sigma$, where σ is a positive scalar. Similarly, suppose there are $\binom{n}{u}$ subtypes of unreliable agents, each with population measure $\binom{n}{r}\sigma$.

Each subtype τ is identified with the set $S_\tau \subset S$ of its bad states (r in number if reliable, u in number if unreliable). All agents t of subtype τ have endowments equal to 1 is $s \in S \backslash S_\tau$, and equal to 0 if $s \in S_\tau$.

The reader can verify that each agent has the right probability of accident (r/n if reliable, u/n if unreliable), and that in *every* state the appropriate fraction of reliable and unreliable agents have accidents.

Turning to our numerical example, $\text{prob}^R(B) = 1/3$ and $\text{prob}^U(B) = 2/3$. Hence $\#S = 3$. There are $\binom{3}{1} = 3$ reliable subtypes, each of measure $\binom{3}{2}\sigma = 3\sigma = 1$ (we set $\sigma = 1/3$), and $\binom{3}{2} = 3$ unreliable subtypes, each of measure $\binom{3}{1}\sigma = 1$.

So let there be $H = 6$ agents, $S = 3$ states of nature, and one good $L = 1$ in each state. Suppose agents have no utility for consumption at $t = 0$, and that they have the same utility at $t = 1$

$$u(x_1, x_2, x_3) = \sum_{s=1}^{3} \log(x_s).$$

The endowments of the agents are

$$e^1 = \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}; \; e^2 = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}; \; e^3 = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix};$$
$$e^4 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}; \; e^5 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}; \; e^6 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}.$$

Notice that the probability of accident for each of the first three "reliable" agents is 1/3, and the probability of accident for each of the last three "unreliable" agents is 2/3. Moreover, in *every* state precisely 1/3 of the reliable type, and 2/3 of the unreliable type, have an accident.

The Arrow–Debreu equilibrium for this economy is $p = (1, 1, 1)$, and $x^h = (2/3, 2/3, 2/3)$ for $h \in \{1, 2, 3\}$, and $x^h = (1/3, 1/3, 1/3)$ for $h \in \{4, 5, 6\}$. This allocation is not achievable via insurance when agent types cannot be observed.

4.3 Insurance with Markets as Intermediaries

Rothschild and Stiglitz introduced oligopolistic companies that provided insurance. Here agents insure each other through the market. They buy pools of promises, and sell promises which allow for idiosyncratic delivery and default.

We take the promises of all assets to be the same:

$$R_j = R_0 \equiv \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \quad \text{for all } j \in J.$$

We take default penalties to be

$$\lambda_{sj}^h = \begin{cases} \infty & \text{if } e_s^h = 1 \\ 0 & \text{if } e_s^h = 0 \end{cases}$$

Notice that the penalty is infinite when agents have the resources to pay, and 0 otherwise. They do not depend on the name of the defaulter, but they do depend on his circumstances. The information required to impose them is identical to the sort of information an insurance company must obtain to verify that an accident has occurred. Indeed we use these penalties precisely in order to render insurance a special case of default.

By combining a long position with a short position on which there may be default, loan (asset) markets may be interpreted as insurance markets. Suppose some asset j has anticipated delivery rates $K_{sj} = K$ for all states s. Consider an agent h who buys and sells one unit of asset j, delivering fully in states s with $e^h_s = 1$ and defaulting completely in states s with $e^h_s = 0$. On net he obtains an "insurance policy" that pays him \$$K$ in every "bad" state s with $e^h_s = 0$, and takes \$$(1 - K)$ in every "good" state with $e^h_s = 1$.

5 Insurance and Adverse Selection without Signalling

Suppose we have a single asset promise $R_0 = (1, 1, 1)$, with $Q^h_0 = \infty \; \forall h \in H$. With only one commodity in each state in period 1, there is no further trade in period 1, and w.l.o.g. we can take $p_s = 1$ for all s. Given the default penalties, all agents will fully deliver in their good states, and fully default in their bad states. We can think of this model as one big insurance contract, with adverse selection. The sellers $h \in \{1, 2, 3\}$ default 1/3 of the time, while the sellers $h \in \{4, 5, 6\}$ default 2/3 of the time. A buyer must anticipate that he may get more sellers of the bad type than of the good type.

In the unique equilibrium derived below, the unreliable agents $h \in \{4, 5, 6\}$ sell and buy twice as much of the asset as the reliable agents $h \in \{1, 2, 3\}$. Hence $K_{s0} = 2/3 \times 1/3 + 1/3 \times 2/3 = 4/9$, $\forall s \in S$. Furthermore, $\theta^h_0 = \varphi^h_0 = 3/5$ for $h \in \{1, 2, 3\}$ and $\theta^h_0 = \varphi^h_0 = \frac{6}{5}$ for $h \in \{4, 5, 6\}$, and $x^1 = (4/15, 2/3, 2/3)$, $x^2 = (2/3, 4/15, 2/3)$, $x^3 = (2/3, 2/3, 4/15)$, and $x^4 = (1/3, 8/15, 8/15)$, $x^5 = (8/15, 1/3, 8/15)$, $x^6 = (8/15, 8/15, 1/3)$. Compared to the Arrow–Debreu equilibrium, reliable agents are doing much worse since their insurance rates are debased by the unreliable agents. The unreliable agents are much better off than in the Arrow–Debreu equilibrium because they benefit from being pooled with the reliable agents.

The reader can verify that the equilibrium is correct by calculating that the marginal utilities to reliable and unreliable agents of buying and selling a unit of the asset is 3.

By selling and buying one unit of the single asset (with fixed delivery rates $K_{s0} = K$), every agent gets K in his bad state, and gives up $1 - K$ in his good state. Since the agents perceive K as fixed, this implicitly defines a price $q = (1 - K)/K$ of consumption in the bad state in terms of the good state. All agents are effectively maximizing $\text{prob}(G) \log x_G + (1 - \text{prob}(G)) \log x_B$ subject to the constraint $x_G + q x_B \leq 1$. With these Cobb–Douglas utilities, agents will always choose $x_G = \text{prob}(G)$, as illustrated in Figure 3. This demonstrates that reliable agents h trade half as much as unreliable agents h', and hence that $K = 4/9$, and hence that $1/q = 4/5$, and hence that $x^h_B = 1/3 \times 1/q = 4/15$, $x^{h'}_B = 2/3 \times 1/q = 8/15$, and so on. In this analysis the price π of the asset played no role, since all agents are buying and selling the same asset. In the next version, with multiple traded assets, the prices π_j are important.

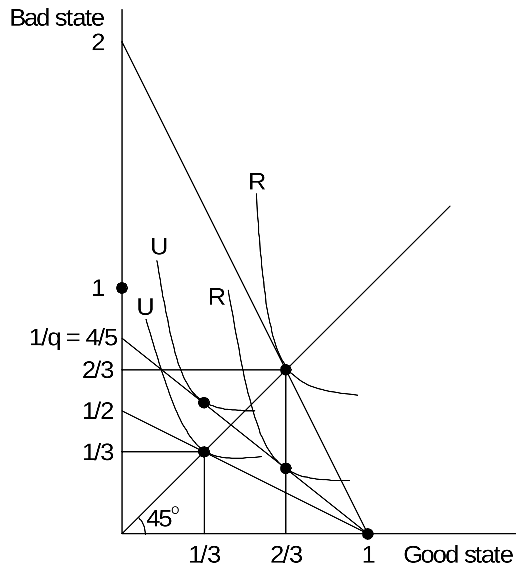

Figure 3. Pooling equilibrium with one asset

Every agent begins with an endowment of 1 in his "good" state(s) and 0 in his bad state(s). The top budget line represents the actuarially fair odds for the reliable agents, for whom the good state is twice as likely as the bad state. At those fair odds (in the Arrow–Debreu equilibrium) they completely insure by moving to the 45° line and consuming 2/3 in every state. Similarly the unreliable agents have a fair odds budget set that reflects the fact that for them the good state is only half as likely as the bad state. In the Arrow–Debreu equilibrium they completely insure by consuming 1/3 in every state.

When the odds are 4:5 of good to bad, the unreliable agents take advantage of the actuarially favorable odds to overinsure, while the reliable agents underinsure because for them the odds are unfair. The odds of 4:5 are closer to 1:2 than to 2:1, reflecting the fact that the unreliable agents take out twice as much insurance. Note finally that the assumption of Cobb–Douglas utilities fixes the same consumption for each agent in his good state, no matter what the budget line.

6 Signalling with Exclusivity Constraints: Separating Insurance

6.1 Signalling Economy with Exclusivity Constraints

Rothschild and Stiglitz (1976) made the important observation that adverse selection in insurance markets might lead to the same kind of inefficient signalling that Spence had earlier postulated would arise in labor markets. In labor markets, Spence (1973) argued that agents with high ability would purchase expensive and unproductive education simply to signal that they were indeed of high ability. In insurance markets, Rothschild and Stiglitz argued, agents would commit themselves exclusively to contracts with low insurance in order to signal that they were reliable. Rothschild and Stiglitz went on to suggest that with signalling there might not be any equilibrium in insurance markets. In their scenario, firms were oligopolistic and could offer insurance contracts at prices visibly different from the market prices prevailing in equilibrium, which had the effect of luring customers away. But in our scenario there is perfect competition: firms do not set the price, the market sets the price. In contrast to Rothschild–Stiglitz, we find that equilibrium always exists in our scenario. But the important point, that signalling can be inefficient, remains intact.

Rothschild and Stiglitz proposed a severe signalling budget set in which agents can sell some contracts which commit them not to sell any other contracts. We can capture this idea by adding to our example additional assets $j = 1, ..., J$ which make the same promises as before, $R_j = (1, 1, 1)$, and with the quantity constraints $Q_j^h = Q_j$, with $Q_j < Q_{j+1}$, for $j = 1, ..., J-1$. By taking the grid size $\equiv \max_{j \in J}\{(Q_{j+1} - Q_j)\}$ to be small, and Q_J and $1/Q_1$ to be large, we can approximate a continuous choice of quantities.

Exclusivity gives the budget sets

$$\underline{B}^h(p, \pi, K) = \{(x, \theta, \varphi, D) \in B^h(p, \pi, K) : \varphi_j > 0 \text{ implies } \varphi_i = 0 \text{ for } i \neq j\}.$$

The asset sales φ must now lie in the set

$$C = \{\varphi \in \mathbb{R}_+^J : \varphi_j \leq Q_j \ \forall j \in J, \ \varphi_j > 0 \Rightarrow \varphi_i = 0 \ \forall i \neq j\}.$$

The set C is star shaped, with each tentacle representing the quantities sellable in a different asset. Clearly $\underline{B}^h$, though nonconvex, satisfies the requirements of Theorem 2.

By symmetry we can suppose $K_{sj} = K_j$ for all $s \in S$. So from the point of view of a buyer, all assets are perfect substitutes. Hence if two assets i and j are traded, their prices must satisfy $\pi_i/K_i = \pi_j/K_j$. Since we can always scale the vector of asset prices arbitrarily, we may assume $\pi_j = K_j$ for all actively traded assets. If an asset is not actively traded, then we must have $\pi_j \geq K_j$. Conversely, if there is at least one actively traded asset, then all agents would be happy not to buy any other asset j with $\pi_j \geq K_j$.

6.2 Price–Quantity Lines

Suppose now that some asset j sells for π_j. An agent who contemplates selling one unit of asset j would receive π_j. He could use this money to buy π_j/π_i units of active asset i, getting deliveries $(\pi_j/\pi_i)K_i = \pi_j$ in every state. The agent has to deliver 1 in all his good states, and nothing otherwise. Thus on net he receives π_j in his bad state and gives up $1 - \pi_j$ in his good state. Implicitly that defines a *price* $q_j = (1 - \pi_j)/\pi_j$, indicating how much consumption must be given up in the good state to get one unit of consumption in the bad state. If a *quantity* φ_j is sold of the asset, then the agent gives up $\varphi_j - \varphi_j \pi_j$ in his good state, and gets $\varphi_j \pi_j$ in his bad state. Thus his consumption must lie at the intersection of the price line and quantity line described in Figure 4 below.

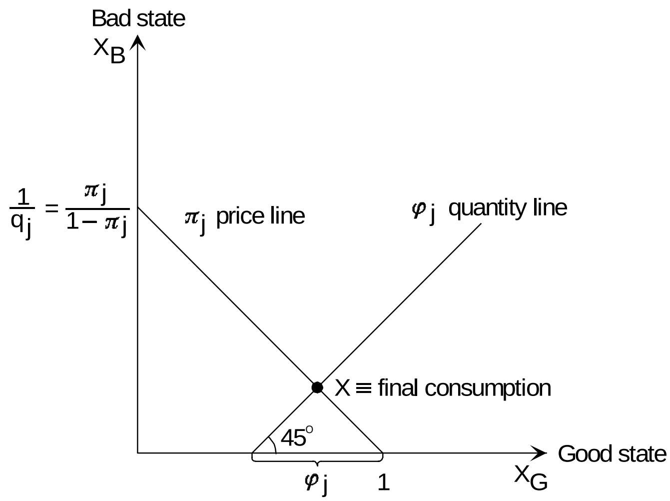

Figure 4. Final consumption from selling *exclusively* φ_j units at price π_j

Each agent is a price taker and a quantity chooser. If the asset has quantity constraint Q_j, then a seller can choose any $\varphi_j \leq Q_j$, giving him the opportunities indicated by the bold interval in Figure 5.

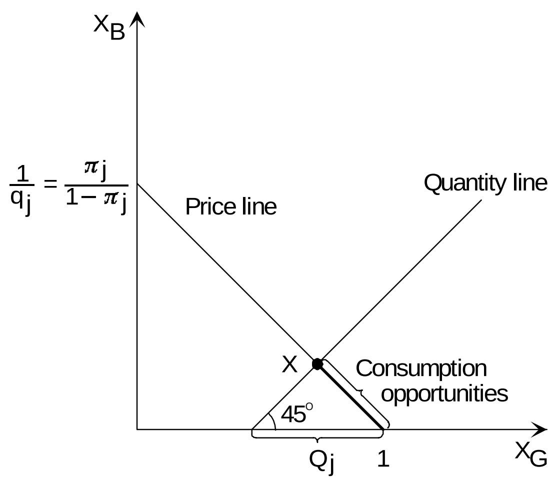

Figure 5

It is important to note that if an agent's indifference curve through X does not intersect the interval of consumption opportunities at any other point, then he will choose $\varphi_j = Q_j$ if he sells asset j.

6.3 Impossibility of Pooling Equilibrium with Exclusivity Constraints

Suppose we are at an equilibrium in which every agent is selling a positive amount of the same asset j^*, with quantity constraint Q_{j^*}. Let final consumption be denoted X^R and X^U, on indifference curves I^R and I^U, respectively. No matter what its price π_{j^*}, we know that the unreliable will sell at least as much as the reliable: $\varphi^U_{j^*} \geq \varphi^R_{j^*}$. Hence $K_{j^*} \leq 1/2 = 1/2 \cdot 1/3 + 1/2 \cdot 2/3$. We shall show that if the grid is fine enough, then we cannot — for all $j \neq j^*$ — find prices π_j at which no agent wants to buy or sell asset j, contradicting that we have an equilibrium in which only asset j^* is traded.

Note that for any quantity limit $0 < Q_j < \varphi^R_{j^*}$, the Q_j quantity line intersects I^R at $y^R(j)$ before it intersects I^U. (See Figure 6.) Clearly $y^R(j)$ lies above the π_{j^*} price line (otherwise R would have chosen $\varphi_{j^*} < \varphi^R_{j^*}$).

Assume that there is a j such that $y^R_G(j) < 1$. Then the points $y^R(j)$ and (1,0) define a price line $\bar{\pi}_j$ indicated by the dotted line.

Assume, furthermore, that $\bar{\pi}_j < 2/3$.

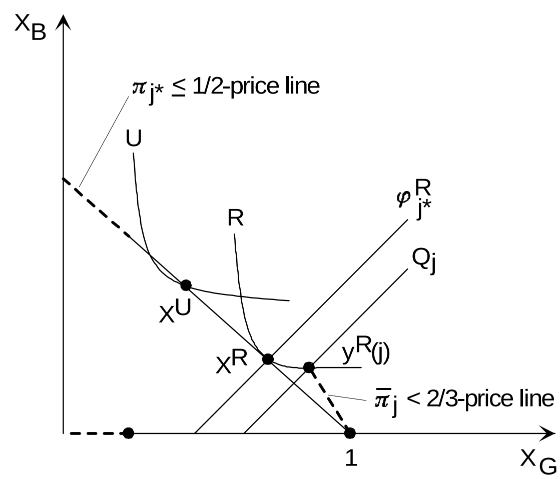

Figure 6

If $\pi_j < \bar{\pi}_j$, all agents strictly prefer not to sell asset j. By the refinement, $K_j = 1$. But then $\pi_j < \bar{\pi}_j < 2/3 < 1 = K_j$, and all agents will rush to buy asset j, contradicting that it is not traded.

On the other hand, if $\pi_j > \bar{\pi}_j$, reliable agents can achieve strictly higher utility selling asset j, contradicting that they chose not to sell it.

Finally, if $\pi_j = \bar{\pi}_j$, then at any perturbation, only the reliable agents could conceivably be induced to sell it. By the refinement $K_j \geq 2/3$. All agents will rush to buy it, since $K_j/\pi_j = K_j/\bar{\pi}_j \geq (2/3)/\bar{\pi}_j > 1$, a contradiction.

We must now justify our two assumptions. They follow trivially from a small grid size, for then taking j with Q_j just barely less than $\varphi^R_{j^*}$, we get that $y^R(j)$ is very near X^R, hence $y^R_G(j) < 1$ and $\bar{\pi}_j$ is near $\pi_{j^*} \leq 1/2 < 2/3$. We must rule out the possibility that $\varphi^R_{j^*} \to 0$ as the grid size goes to zero.

We argue that $\varphi_{j^*}^R \geq \bar{\varphi} > 0$, where $\bar{\varphi}$ is given in Figure 7, independent of the grid size. Since in equilibrium all assets must have $K_j \geq 1/3$, we may suppose $\pi_0 \geq 1/3$. At worst, the reliable agent must achieve the utility he could get from selling asset 0 at price 1/3. Since $u'(0) = \infty$, this would indeed give him more utility than his endowment (1,0). Hence he would never agree to sell less than $\bar{\varphi} > 0$ of any asset priced at $\pi_j = K_j = 1/2$. A fortiori, for any price $1/3 \leq \pi_j \leq 1/2$, he would insist on an even higher minimum sale of asset j.

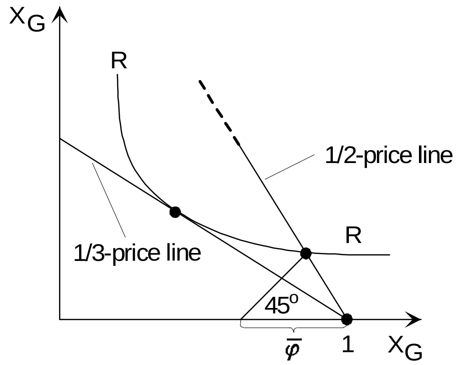

Figure 7

Our equilibrium refinement thus captures the same spirit as the Rothschild–Stiglitz definition of no entry equilibrium in eliminating the pooling equilibrium.

6.4 Existence of Separating Equilibrium

The only equilibrium Rothschild and Stiglitz found is the "separating" equilibrium in which each type of agent sells a different asset. Rothschild and Stiglitz observed that in such an equilibrium the unreliable types should feel unconstrained by the quantity restriction while the reliable types should feel quantity constrained. Moreover, the unreliable types should be indifferent to either of the two contracts, while the reliable types should strictly prefer their quantity constrained contract. Indeed, once we impose the Rothschild–Stiglitz exclusivity restriction, we get this sort of equilibrium, though not quite exactly because our menu of quantity constraints is finite. The difference is that our separating equilibrium always exists, whereas Rothschild–Stiglitz found robust regions of nonexistence.

We claim that with the exclusivity condition ($\varphi_j^h > 0$ for at most one contract j), there is essentially a unique equilibrium in terms of consumption. We first describe its qualitative features before computing it for our numerical example. See Figure 8.

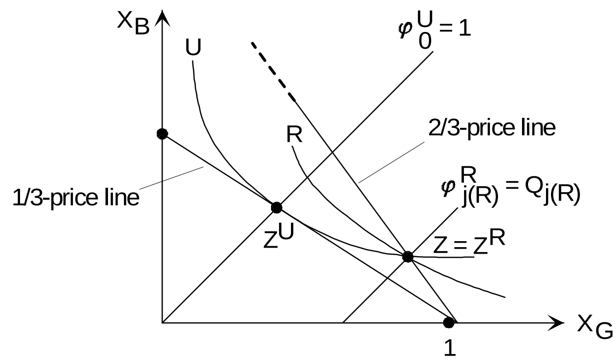

Figure 8

Let the unreliable agents trade (i.e., buy and sell) one unit of asset 0, with $\pi_0 = K_0 = 1/3$, to obtain their optimal consumption Z^U on the 1/3-price line. Let Z be the intersection of the U-indifference curve through Z^U with the 2/3-price line. *Assume* that there exists an asset $j(R)$ such that Z lies on the $Q_{j(R)}$-quantity line. Then let the reliable agents trade $\varphi^R_{j(R)} = Q_{j(R)}$ units of asset $j(R)$, with $\pi_{j(R)} = K_{j(R)} = 2/3$, to obtain the consumption $Z^R = Z$.

To check that we have a genuine equilibrium, we must price all the untraded assets in a manner that satisfies our equilibrium refinement.

If $1 \leq Q_j$, set $\pi_j = K_j = 1/3$.

If $Q_{j(R)} < Q_j < 1$, then the Q_j-quantity line intersects the U-curve (through Z^U) at $\tilde{Z}^U$, before it intersects the R-curve (through Z^R). Set $\pi_j = K_j$ in accordance with the dotted line in Figure 9, which connects (1,0) to $\tilde{Z}^U$.

If $Q_j < Q_{j(R)}$, then the Q_j-quantity line intersects the R-curve at $\tilde{Z}^R$ before the U-curve. If $\tilde{Z}^R_G < 1$, set $\pi_j = K_j$ in accordance with the other dotted line in Figure 9, which connects (1,0) to $\tilde{Z}^R$.

If $Q_j < Q_{j(R)}$ and $\tilde{Z}^R_G \geq 1$, set $\pi_j = K_j = 1$.

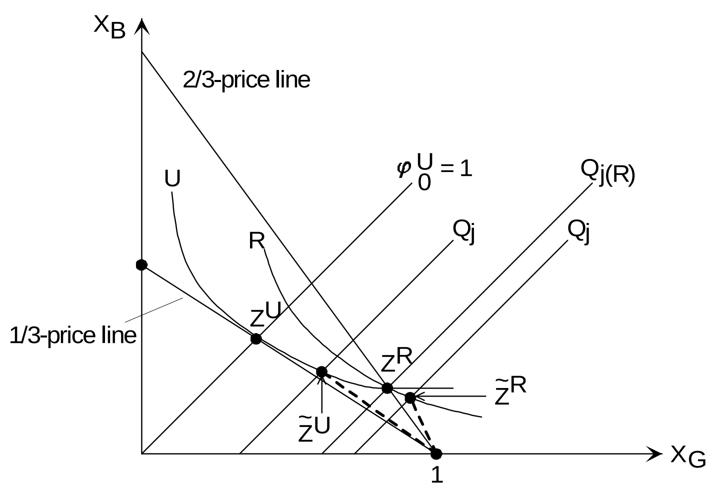

Figure 9

For every $j \notin \{0, j(R)\}$ we show that no agent can strictly improve by trading j.

When $1 < Q_j$ (and so $\pi_j = 1/3$), the unreliable agents are indifferent to trading (1 unit of) j and (1 unit of) asset 0, while the reliable types are strictly worse off trading j.

When $Q_{j(R)} < Q_j \leq 1$, the unreliable agents are indifferent to trading (Q_j units of) j and (1 unit of) asset 0, while reliable types are strictly worse off trading j.

When $Q_j < Q_{j(R)}$ and $\pi_j < 1$, the reliable agents are indifferent to trading (Q_j units of) j and ($Q_{j(R)}$ units of) asset $j(R)$, while the unreliable agents are strictly worse off trading j.

Finally, when $Q_j < Q_{j(R)}$ and $\pi_j = 1$, no agent can do better switching from j^* to j. (See Figure 10.)

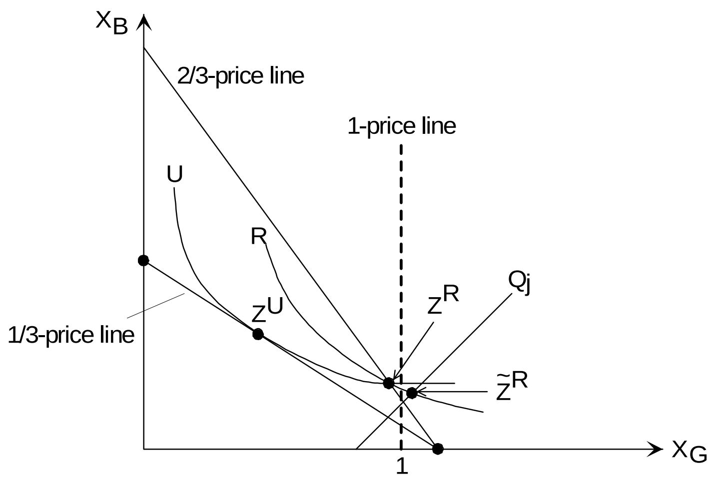

Figure 10

The 1-price line is the vertical line above (1,0). Since the R-indifference curve is strictly downward sloping, we have $\tilde{Z}^R_G \geq 1$ if and only if the Q_j quantity line intersects the 1-price line on or below the R-indiffrence curve. Thus the consumption opportunities (from (1,0) to this intersection point) cannot make the reliable agent better off than his equilibrium consumption Z^R. The unreliable agents are strictly better off at Z^U.

In every case $j \notin \{0, j(R)\}$, with $\pi_j < 1$, either the reliable types or the unreliable types (but never both) are indifferent to switching from their equilibrium trade to j. In each of these cases, $\pi_j = K_j$ is at least as high as the delivery rates of the only type willing to switch there.

We check now that these prices satisfy our equilibrium refinement. Given any $\varepsilon > 0$, choose the macro variables $(p(\varepsilon), \pi(\varepsilon), K(\varepsilon))$ in $E(\varepsilon)$ to be *identical* to the macrovariables (p, π, K) for E. For assets j with $\pi_j = 1$, we have $K_j(\varepsilon) = K_j = 1$ and so the refinement condition is automatically satisfied. For each asset j with $\pi_j < 1$, let a small measure (less than ε) of the relevant type switch to trading asset j. They are still optimizing, and all markets continue to clear. Moreover, $K_j(\varepsilon) = K_j$ is no less than the realized delivery rate on j in $E(\varepsilon)$. Since boosting of delivery rates is permitted in $E(\varepsilon)$ for assets that are untraded in E, we have constructed a valid perturbation.

We now turn to the general situation in which there is no asset j such that the Q_j quantity line contains Z. (See Figure 11.)

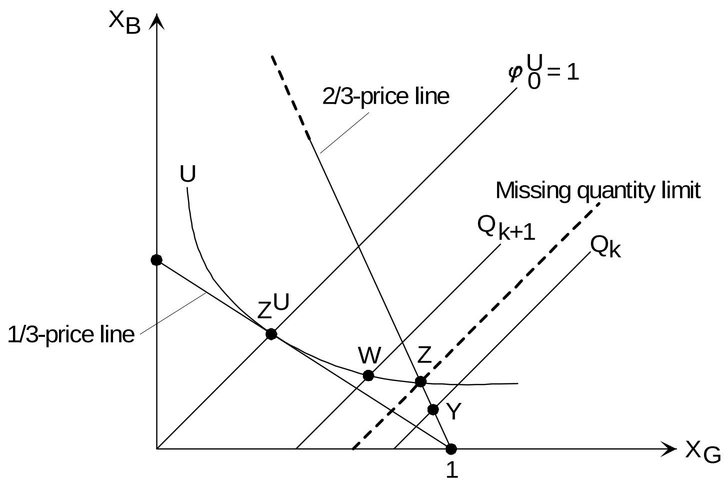

Figure 11

Let Q_{k+1}, Q_k be consecutive quantities in the grid which “trap” the missing quantity in between. Denote

$W \equiv$ intersection of the U-indifference curve (through Z^U) with the Q_{k+1}-quantity line

$Y \equiv$ intersection of the 2/3-price line with the Q_k-quantity line.

Case 1 (Figure 12) The reliable type (weakly) prefer Y to W.

Then define $Z^R = Y$ and $j(R) = k$, and proceed exactly as before to price the untraded assets and to construct the perturbation.

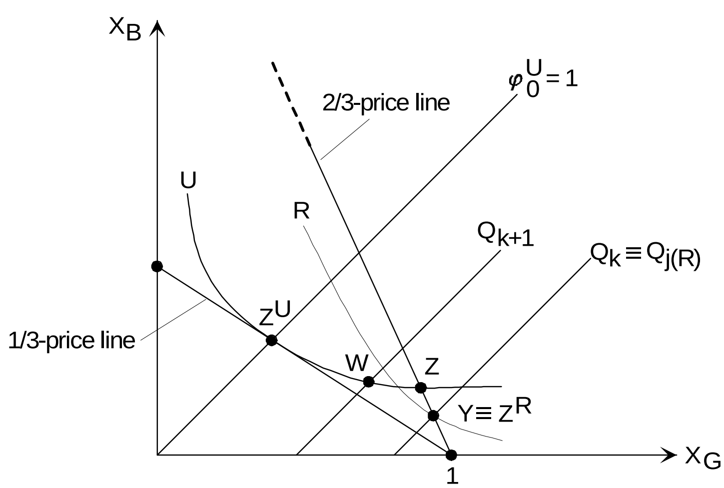

Figure 12

6.5 Nearly Separating Equilibrium

Case 2 (Figure 13) The reliable prefer W to Y.

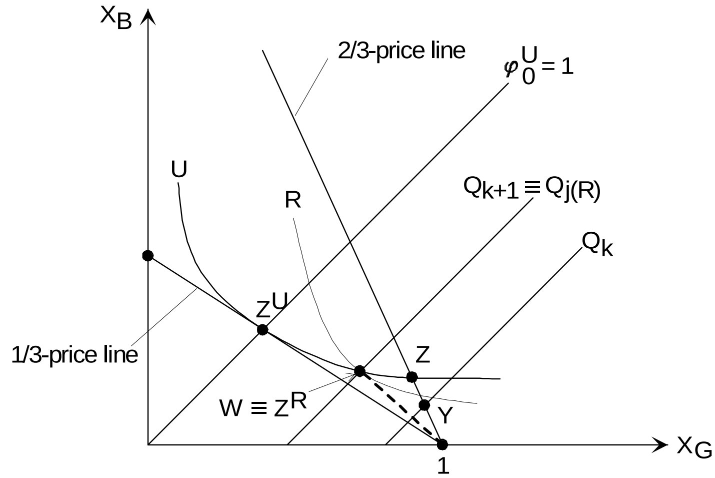

Figure 13

In this case we do not get a pure separating equilibrium, but an equilibrium with a slight degree of "mixing." Let all the reliable agents trade $Q_{k+1} \equiv Q_{j(R)}$ units of asset $k+1 \equiv j(R)$ and consume $W \equiv Z^R$. Set $\pi_{j(R)} = K_{j(R)}$, in accordance with the dotted line joining (1,0) to W in Figure 13. The new feature of this equilibrium is that some unreliable agents *also* trade the asset $j(R)$. In fact, just enough of them trade $j(R)$ so that the delivery rate falls from 2/3 to $K_{j(R)}$. The rest of the U population acts as before, trading $\varphi_0 = 1$ units of asset 0.

Notice that the pricing of untraded assets and the perturbation work exactly as before. Also notice that the degree of mixing goes to zero with the grid size (since then $W \to Z$ and $\pi^* \to 2/3$).

Notice also that the slight mixing has absolutely nothing to do with the population proportion of reliable and unreliable types, and so nothing to do with the nonexistence in Rothschild–Stiglitz. Had we taken all possible Q, we would always obtain the separating equilibrium.

6.6 The Numerical Example

In our numerical example, let us add assets $j = 1, ..., 100$ to asset 0, with $Q_j = j/30$ (and so, with grid size 1/30). The reader can check that this puts us in Case 1 with $j(R) = 9$. See Figure 14.

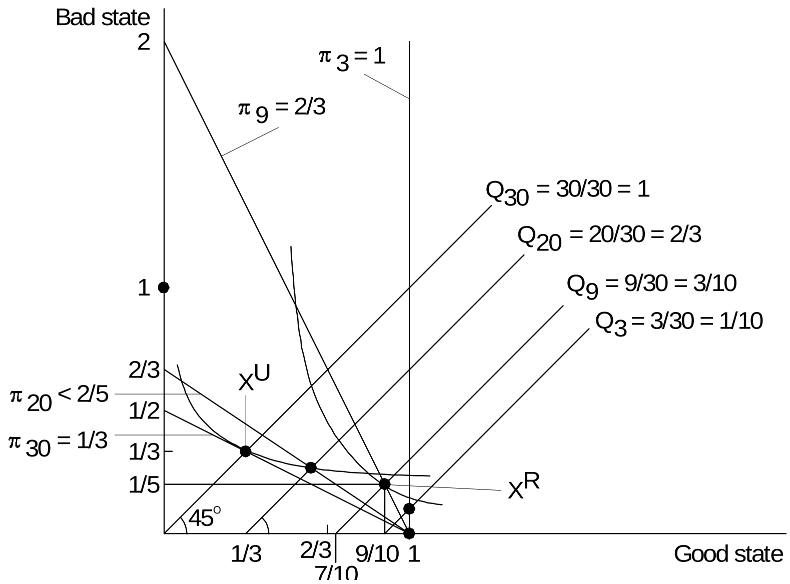

Figure 14

It is worth pointing out that we were able to exploit the special structure of the insurance economy (like the single crossing property) to construct the equilibrium, and to verify it to be so. In general, constructing equilibrium can be quite difficult, but there is no question of its existence. Theorem 2 assures us that equilibrium exists even without special structure, for example, even if reliable agents have utlity $v \neq u$ so that v and u indifference curves cross more than once.

Rothschild and Stiglitz correctly noted that the separating equilibrium allocation and price system is well-defined and feasible independent of the proportion of reliable agents. By contrast, observe that the pooling equilibrium improves in utility terms as the proportion of reliable agents converges to 1, eventually Pareto dominating the separating equilibrium. Rothschild and Stiglitz went on to claim that if nearly all the agents are reliable, then the separating equilibrium could not be an equilibrium, because some contract such as asset 30 with its more generous constraint $Q_{30} = 1$ would break the equilibrium. Their paper is not precise about how expectations are formed when assets are not traded, but the idea is that if it was expected that the sellers of asset 30 were in the same proportion as the population as a whole, then in a population consisting almost entirely of reliable agents, the corresponding K_{30} would be nearly 2/3 and the price it would fetch would be $\pi_{30} = 2/3$. This price (and its generous quantity constraint) definitely would lure away sellers of both types and, so Rothschild and Stiglitz argue, justify the expectations $K_{30} = 2/3$, and thus upset the separating "equilibrium."

However, such an expectation is hasty, since agents do not all have the same incentive to switch to the new contract. We have set the price in equilibrium of asset 30 at $\pi_{30} = 1/3$. At this price the unreliable agents are just indifferent to switching from asset 0 into asset 30, whereas the reliable agents are not close to wanting to sell asset 30. Even if some agent offered to buy asset 30 at a price of 1/2, only unreliable agents would rush to sell it. Not until the price reaches .53 would the reliable agents become interested in selling asset 30. Thus we feel justified in setting the expectations of delivery for asset 30 at $K_{30} = 1/3$, and the price at $\pi_{30} = 1/3$. If an agent did for some remarkable reason offer to pay 2/3 for asset 30, he would be deluged with offers from sellers, so many in fact that he could never accommodate them all. A natural reaction would be to lower his buying price until the number of sellers fell. But as we just saw, to reduce demand sufficiently he would end up selling only to the unreliable types, as we have presumed.

We are in agreement with the concern of Rothschild and Stiglitz about the separating equilibrium when there is a high proportion of reliable agents. But the problem is not the nonexistence of equilibrium. (We constructed an equilibrium in our example, and in general, Theorem 2 guarantees that equilibrium exists.) The problem is its inefficiency.

References

[1] Akerlof, G., 1970. "The Market for Lemons: Qualitative Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84: 488–500.

[2] Allen, F. and D. Gale, 1988. "Optimal Security Design," *Review of Financial Studies*, 1: 229–263.

[3] Cho, I. K. and D. Kreps, 1987. "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102: 179–221.

[4] Dubey, P., J. Geanakoplos and M. Shubik, 2000. "Default in a General Equilibrium Model with Incomplete Markets," Cowles Foundation Discussion Paper, No. 1247.

[5] Dubey, P., J. Geanakoplos and M. Shubik, 2001. "Default and Punishment in General Equilibrium with Incomplete Markets," Cowles Foundation Discussion Paper No. 1304.

[6] Dubey, P. and J. Geanakoplos, 2001. "Insurance and Signalling in Perfect Competition," forthcoming Cowles Foundation Discussion Paper.

[7] Gale, D., 1992. "A Walrasian Theory of Markets with Adverse Selection," *Review of Economic Studies*, 59: 229–255.

[8] Helpman, Elhanan and Jean-Jacques Laffont. 1975. "On Moral Hazard in General Equilibrium Theory," *Journal of Economic Theory*, 10: 8–23.

[9] Pesendorfer, W., 1995. "Financial Innovation in a General Equilibrium Model," *Journal of Economic Theory*, 65: 79–116.

[10] Prescott, E. and R. Townsend, 1984. "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard," *Econometrica*, 52: 21–45.

[11] Riley, J., 1979. "Informational Equilibrium," *Econometrica*, 47: 331–359.

[12] Rothschild, M. and J. Stiglitz, 1976. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 90: 629–650.

[13] Spence, M., 1973. "Job Market Signalling," *Quarterly Journal of Economics*, 87: 355–374.

[14] Stiglitz, J. and A. Weiss, 1981. "Credit Rationing in Markets with Imperfect Information," *American Economic Review*, 72: 393–410.

[15] Wilson, C., 1977. "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory*, 16: 167–207.