

**COMPETITIVE POOLING: ROTHSCHILD-STIGLITZ
RECONSIDERED**

BY

**PRADEEP DUBEY
AND
JOHN GEANAKOPLOS**

COWLES FOUNDATION PAPER NO. 1048



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY**

**Box 208281
New Haven, Connecticut 06520-8281
2003**

COMPETITIVE POOLING: ROTHSCHILD-STIGLITZ RECONSIDERED

PRADEEP DUBEY AND JOHN GEANAKOPOLOS

We build a model of competitive pooling, which incorporates adverse selection and signaling into general equilibrium. Pools are characterized by their quantity limits on contributions. Households signal their reliability by choosing which pool to join. In equilibrium, pools with lower quantity limits sell for a higher price, even though each household's deliveries are the same at all pools. The Rothschild-Stiglitz model of insurance is included as a special case. We show that by recasting their hybrid oligopolistic-competitive story in our perfectly competitive framework, their separating equilibrium always exists (even when they say it does not) and is unique.

I. INTRODUCTION

Traditional general equilibrium theory treated insurance as a special case of securities with contingent payoffs. A household with low endowment in some state could "insure" himself by buying a security which delivered when he most needed the money.

What is missing from this traditional approach is pooling. In practice, an insurance company issues a generic contract, to pay in case of "accident." Different clients sign the same insurance contract, but they thereby purchase different securities, because their "accident" states are different. The shareholder in the insurance company in effect holds a pool of different liabilities. Pooling inevitably leads to adverse selection because households with more probable accident states have incentive to take out more insurance and therefore tend to be more than proportionately represented in the pool.

In this paper we develop a simple theory of competitive pooling, starting from the point at which traditional general equilibrium theory stopped. Pooling encompasses many examples in addition to insurance. A primitive example is the land pool in which households contribute part of their land to a common pool (or cooperative), whose collective output is then distributed in proportion to the number of acres put in. Since the output of different land differs, the shareholders of the cooperative receive the average output and not the output of any single contributor. A more modern example is provided by the huge mortgage pools traded on Wall Street. Different homeowners essentially sell the

same generic promise, to deliver \$1 a month for 30 years, but because of idiosyncratic prepayments and defaults, their actual deliveries are quite different. Shareholders in these pools again receive the average of the homeowner deliveries.

In our model, an agent h chooses to sell φ_j^h promises into pool j , and is obliged to deliver an exogenously prescribed d_j^h per promise. Different sellers deliver differently, but their promises cannot be distinguished. We suppose that buyers and sellers do not trade bilaterally, but through the anonymous pool. The buyers (shareholders) of pool j receive a pro rata share of all its different sellers' deliveries. Each share of pool j delivers

$$K_j = \frac{\sum_h \varphi_j^h d_j^h}{\sum_h \varphi_j^h}.$$

The shareholder of pool j does not know, or need to know, the identities of the sellers or the quantities of their sales. All that matters to him is the price π_j of the share and the delivery rate K_j .

Pooling dramatically reduces the information needed to buy a diversified portfolio of risks: instead of forecasting individual deliveries K_j^h for many different individuals h , a buyer need only concern himself with a single average delivery K_j . Figuring out K_j^h for one individual is typically no less difficult than estimating K_j for a pool with a large population. Thus, pooling overcomes the costly information processing problems inherent in multiple bilateral negotiations, and is one reason why it is becoming so prevalent in modern economies.

The pooling also leads to *adverse selection*, since a buyer must worry that unreliable sellers with a proclivity for lower deliveries will tend to sell more promises into the pool, worsening the anticipated rate K_j . *Signaling*, by publicly committing oneself to a small quantity of sales, therefore has an important role to play, because it suggests to the buyer that deliveries may be more reliable. To incorporate it in our model, we suppose that there are many pools j , each with its own quantity limit Q_j imposed on sales into the pool. No one is permitted to sell into more than one pool. This opens up the opportunity for agents to signal their restraint by selling into a pool with low Q_j .

By enabling each agent to trade anonymously as part of a large aggregate, pooling already takes us part of the way toward perfect competition. We fully get there by postulating that all

agents view (Q_j, π_j, K_j) at each pool j as fixed. Perfect competition thus further reduces the information a buyer requires: there is no need for him to forecast how the delivery rate at any pool would vary if the price were changed, since he cannot change the price.

The terms Q_j of pool j are set exogenously, just as the location, date, and quality of a commodity are in traditional general equilibrium theory. The prices π_j , the anticipated delivery rates K_j , and the trades at each pool j are all determined endogenously at equilibrium by the market forces of supply and demand.¹ It will turn out that most pools have no trade in equilibrium, and thus the quantities Q_j of *active* pools may also be thought of as endogenous.

Adverse selection and signaling have been recognized for a long time by the burgeoning field of contract theory, in which a classic problem is how insurance companies will design contracts to protect themselves from adverse selection. Rothschild and Stiglitz [1976] wrote a pioneering article in this field, in which oligopolistic, risk-neutral insurance firms design a menu of contracts for a continuum of private agents in face-to-face meetings. Each contract j specifies precisely how much insurance Q_j to take out and at what rate π_j , and clients are prohibited from choosing more than one contract. Not only the terms Q_j , but also the prices π_j , are viewed as part of the contract. The insurance companies are imagined to deduce the change in reliability K_j of the clientele that would be forthcoming as they vary either Q_j or π_j . Rothschild and Stiglitz found that when equilibrium exists, only two contracts are actually offered, with reliable agents choosing one and unreliable agents choosing the other. They noted, however, that there are robust regions in which no equilibrium, as they defined it, exists.

We recast this story in our perfectly competitive setting, retaining only the continuum of agents. We do not have gargantuan, risk-neutral insurance companies—we have pools. Diverse groups of small risk-averse households trade promises through these anonymous pools. Since the assets bought are pools of promises, and those sold permit idiosyncratic deliveries, the net effect is that households insure each other through the pools. *Every* agent is a price taker. Yet the model is subtle enough to unambiguously determine which insurance contracts will emerge

1. In the simple model of pooling we describe, it is always the case that $\pi_j = K_j$ in equilibrium.

as actively traded. It is not the managers of oligopolistic firms, but the invisible hand of perfect competition that takes over the role of designing contracts.

By recasting the hybrid oligopolistic-competitive story of Rothschild-Stiglitz in our perfectly competitive framework, we simplify the analysis and obtain stronger conclusions. Equilibrium *always exists* in our model (even when they say it does not) and is *always unique*.

In our equilibrium adverse selection plays a prominent role: two pools i and j with identical deliveries $d_i^h = d_j^h$ for all households h nevertheless sell for different prices $K_i = \pi_i > \pi_j = K_j$ if $Q_i < Q_j$. For all but two pools there is no trade at all, even though each pool Q is open and ready for business at its equilibrium $\pi = K$. The two active pools correspond exactly to the Rothschild-Stiglitz separating equilibrium.

Perfectly competitive pooling conforms to much of the real world. A prominent example is the securitized pass-through mortgage market, in which many homeowners make the same promise but deliver differently. This example is of particular importance to us because it illustrates the difference between competitive and oligopolistic lending. Every mortgage is issued through a bank in a one-on-one meeting between lender and borrower. The contract theory literature emphasizes this relationship. But modern developments have dramatically changed the situation. The banks are allowed to immediately sell the mortgages into a pool, managed by some agency, which then sells shares. The lenders are really the shareholders in the pool, who recognize that they are so small that no matter how much they buy, they will not affect the mortgage rates that the homeowners pay or the delivery rates they receive. The banks do not for the most part have any discretion over whether to make the loans or not. The pooling agency prescribes the criteria for lending, and the bank simply verifies these. The economic analysis thus properly shifts from the game theoretic level of lender and borrower in face-to-face contact, to the pool level of perfectly competitive, anonymous shareholders and borrowers. This anonymity does not rule out adverse selection or signaling. Indeed, the heart of mortgage investing is to understand whether a pool consists mostly of reliable or unreliable agents. Mortgage pools differ in the maximum loan Q that can be taken out. As a rule, pools with smaller loans Q fetch better prices, precisely because they signal that the borrowers are more reliable (they default less, and are less so-

phisticated in prepaying). As a result, a homeowner wishing to get a \$200,000 mortgage will get a better interest rate than a homeowner wishing to take out a \$400,000 mortgage. These mortgage pools total approximately \$3 trillion, so that it appears more and more today that a substantial part, if not yet the majority, of lending is done through pools. Credit card pools, incidentally, constitute another very similar example.

Our treatment is firmly in the tradition of perfect competition, but with one significant twist. When there is no trade in a pool j , potential investors are unable to validate their anticipated K_j with realized deliveries: K_j is not defined by our formula when total sales $\sum_h \phi_j^h = 0$. If K_j were allowed to be arbitrary, it could be so low that no agent would have incentive to join pool j . Thus, agents' whimsical pessimism could render any ad hoc set of pools inactive. To overcome this problem, we consider a simple equilibrium refinement. The idea is to introduce a fictitious seller (say, the government) who contributes an infinitesimal ϵ promise to each pool and makes ultrareliable deliveries $\epsilon d > \epsilon \max_h d_j^h$. This fixes anticipations at the most optimistic level that is consistent (as we shall see) with cautious rationality. In spite of this refinement, many pools will be inactive at equilibrium in our model, but their selection will no longer be ad hoc, and indeed will be unique in the context of insurance that we focus on. (In particular, the refined equilibrium will not depend on the precise deliveries d of the external agent, so long as d is ultrareliable.)

One 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 inactive pools. We argue in Section VIII that our cautious optimism is natural when there are many buyers and sellers in perfect competition. By contrast, the expectations attributed to agents by Rothschild and Stiglitz are not compatible (to our way of thinking) with perfect competition.

In the Rothschild-Stiglitz model, equilibrium is required to be immune to entry by new insurance companies who might offer a contract (Q_j, π_j) that would turn a profit by luring households away from their old contracts. One might well ask whether our equilibrium is immune to entry. The answer is that whatever new Q could be imagined is already present and embodied by one of the pools j , and its associated quantity limit $Q_j = Q$. Its price $\pi_j = K_j = K(Q_j)$ is set by the market. This brings us to the second crucial difference between the Rothschild-Stiglitz model and

ours. For them the terms Q of a contract and the price π enter symmetrically. Thus, they must consider the potential trade at all pairs (Q, π) . For us, the terms Q are given exogenously, and the prices $\pi = K$ are determined by the market. Hence we need consider only one K for each Q . In effect they must consider all contracts in a square, and we need only look at contracts along the diagonal. That is one reason our equilibrium is simpler, and why it always exists. But the remaining contracts are still numerous enough to capture all the relevant economics of adverse selection and signaling.

Perfect competition not only simplifies the equilibrium, but also its refinement. In the contract theory literature, when two parties are in face-to-face meetings, an extensive form game is created, in which the refinements are vastly more complex. They require agents to engage in a long chain of hypothetical reasoning about each other. For example, in the refinement of Cho and Kreps [1987], h must think about what j thinks about what every other player k (including h himself) is thinking about, in order to deduce whether j will be able to deduce who he is dealing with. It presupposes common knowledge of private, individual characteristics; and calls upon each agent not only to think through many iterations, but to believe that others are doing likewise. Our refinement strains credulity less. There is no hypothetical reasoning and no chain. Agents think only about the observable macro aggregates K_j . The concrete, infinitesimal actions of the external agent are relevant only through their impact on the K_j ; indeed, their purpose is to render the K_j observable.

In the modern world one sees many examples of pools of promises, e.g., insurance pools, mortgage pools, credit card pools, and so on. Often entry into a pool is signified by a virtual promise which is identical across agents. It is understood, however, that different households will actually deliver differently. The mechanisms by which these different deliveries come about involve options and default (and give rise to moral hazard). But as long as actual deliveries are foreseen, the analysis we do in this paper will remain relevant in the study of equilibrium. In our model here we take a simpler and more abstract approach in that the deliveries d^h are exogenous, though they depend on the individual h . By enlarging the model to include more fundamental motivations, we could derive the d^h from different individual incentives. For example, in an earlier paper [Dubey, Geanakoplos, and Shubik 2001] we showed how idiosyncratic default

penalties for failure to deliver would lead to different deliveries d^h that could be pooled. In Section IV we discuss several other contexts in which the d^h emerge from underlying microeconomic considerations.

In this paper we prohibit a household from taking out more than one insurance contract (i.e., from selling into more than one pool). This is done to bring our analysis in line with that of Rothschild-Stiglitz. But of course one is then led to ask what would happen if households were free to taken out multiple insurance contracts. This is the topic of our sequel paper [Dubey and Geanakoplos 2001].

II. THE POOLING OF PROMISES

Imagine households $h \in \mathcal{H} = \{1, \dots, H\}$, each of whom has a risky endowment $e_s^h \in \mathbb{R}_+$ (of money), depending on the state of nature $s \in S = \{1, \dots, S\}$. We assume that $\sum_{h \in \mathcal{H}} e^h \gg 0$.

Every household h is risk-averse, and his ex ante utility (for money) is given by a continuous, strictly monotonic, strictly concave function,

$$u^h : \mathbb{R}_+^S \rightarrow \mathbb{R}.$$

We suppose that trade takes place only through contributions to a pool. Later we shall allow households to choose between competing pools.

Each household $h \in \mathcal{H}$ is entitled to contribute $0 \leq \varphi^h \leq Q$ promises to the pool, which oblige him to make state-contingent deliveries $d^h \in \mathbb{R}_+^S$ per unit of promise. The promises $\varphi \equiv (\varphi^1, \dots, \varphi^H)$ are aggregated in the pool and yield an average delivery,

$$K_s(\varphi) = \left(1 / \sum_{h \in \mathcal{H}} \varphi^h \right) \sum_{h \in \mathcal{H}} \varphi^h d_s^h,$$

in every state $s \in S$. (We take $K_s(\varphi) = \text{arbitrary}$, if $\sum_{h \in \mathcal{H}} \varphi^h = 0$.) Household h , who holds φ^h shares of the pool by virtue of contributing φ^h units of his promise to it, ends up therefore with the final bundle $\chi(\varphi^h, K(\varphi)) \in \mathbb{R}^S$, with components,

$$\chi_s(\varphi^h, K(\varphi)) = e_s^h + \varphi^h(K_s(\varphi) - d_s^h),$$

for $s \in S$. The feasible set of contributions available to h is given by

$$\{\theta^h \in [0, Q] : \chi^h(\theta^h, K(\varphi|\theta^h)) \in \mathbb{R}_+^S\},$$

where we have denoted $(\varphi|\theta^h) \equiv (\varphi^1, \dots, \varphi^{h-1}, \theta^h, \varphi^{h+1}, \dots, \varphi^H)$. The rules of the cooperative pool thus define a noncooperative (generalized) game with payoffs $u^h(\chi(\varphi^h, K(\varphi)))$, $h \in \mathcal{H}$.

The reader will notice that we have not separated sales into the pool from the purchase of shares of the pool. Indeed, we have supposed that the sales themselves determine the shares. This simple rule describes how cooperatives function. Of course, modern financial institutions decouple buying and selling. We could easily accommodate this, and our entire analysis would remain intact. But there would be a cost of added notation. Our simple pooling captures the essential idea we mean to exposit, and is sufficient to represent the Rothschild-Stiglitz model of insurance.

III. THE PERFECTLY COMPETITIVE COOPERATIVE

In the game, households must anticipate that their contributions alter the pool quality $K(\varphi)$. When the number of households is very large, this quality effect becomes almost negligible. By ignoring it, any one household can concentrate on the far simpler problem of determining how much of the "net trade" ($K - d^h$) he wishes to acquire.

We now postulate a world in which it is perfectly rational for each household to take K as given, independent of his action. This simplifies the analysis of equilibrium, without compromising the economic phenomena of adverse selection and signaling.

Let us imagine a continuum of households $t \in (0, H]$, where all $t \in (h - 1, h]$ are of type h and are identical: $d^t = d^h, e^t = e^h, u^t = u^h$. Given a measurable choice of actions $\varphi : (0, H] \rightarrow [0, Q]$ (which we also write $\varphi \in [0, Q]^{(0, H]}$), the pool holds $\bar{\varphi} \equiv \int_0^H \varphi^t dt$ units of aggregate promise and delivers $K_s(\varphi) \equiv 1/\bar{\varphi} \int_0^H \varphi^t d_s^t dt$ per unit, if $\bar{\varphi} > 0$. It is clear that no single household in the continuum $(0, H]$ can affect $K_s(\varphi)$ by changing his actions. From this point of view, the trading opportunities are specified by the fixed vector $K = K(\varphi)$. Household $t \in (h - 1, h]$ consumes

$$x_s^t = \chi_s^t(\varphi^t, K) \equiv e_s^h + \varphi^t(K_s - d_s^h)$$

money in each state $s \in S$. His budget set is given by

$$\Sigma^t(K) = \{(\theta, y) \in [0, Q] \times \mathbb{R}_+^S : y = \chi^t(\theta, K)\}.$$

We will say that $(K, \varphi, x) \in \mathbb{R}_+^S \times [0, Q]^{(0, H]} \times [\mathbb{R}_+^S]^{(0, H]}$ is

an equilibrium for the one-pool economy $((u^h, e^h, d^h)_{h \in \mathcal{H}}, Q)$ iff φ and x are measurable and

$$(1) \quad K = (1/\bar{\varphi}) \int_0^H \varphi^t d^t dt \text{ if } \bar{\varphi} > 0$$

$$(2) \quad (\varphi^t, x^t) \in \arg \max_{(\theta, y) \in \Sigma^t(K)} u^t(y) \text{ for almost all } t.$$

Notice that we are silent on how K should be formed when $\bar{\varphi} = 0$. By taking $K = 0$ (or sufficiently small, provided the marginal utilities of u^h are bounded), we can always sustain an *inactive* equilibrium (K, φ, x) in which $\varphi^t = 0$ almost everywhere. With only one cooperative this is not a serious matter, since we lose little by confining our attention to equilibria (K, φ, x) which are *active*, in the sense that $\bar{\varphi} > 0$. But when we consider multiple cooperatives, we will find that many of them must be inactive in equilibrium, and then the choice of their K becomes a crucial issue. By its presence, K "opens" the inactive cooperative's doors for business: every household t knows that he will receive θK in exchange for θd^t . If the cooperative pool is inactive in equilibrium, it is in spite of this trading opportunity, since all households are choosing voluntarily not to go there. But simply having the doors open at every pool is not enough to define a reasonable equilibrium. The *levels* K must be appropriately pinned down, which is the purpose of our refinement in Section VIII.

IV. SOME MORE EXAMPLES OF POOLING

The easiest cooperative to describe is the land pool in which farmers contribute as much of their land as they wish. One might imagine that land put into the cooperative is painted blue, while the land held back is painted red. At a glance the farmers can survey the aggregate blue land held by the cooperative. Farmers, even with especially productive land, may be willing to contribute to the cooperative because in that way they insure themselves against states in which their crop fails relative to the average acre in the pool.

The modern corporation is like a pool. General Electric, for example, has many different businesses, spread all over the world, ranging from financial services to dishwashers. All these subsidiary businesses contribute their profits to the pool owned by shareholders of GE stock.

Asymmetric information provides another reason agents

might deliver differently into the same pool. Suppose that agents make the same state contingent promises R_s , $s \in S$. But suppose that each agent h has a partition Π^h of the states of nature reflecting his private information. When state s occurs, he cannot distinguish it from other states $\sigma \in \Pi^h(s)$ and, if default is prohibited, is forced to deliver

$$d_s^h = \max \{R_\sigma : \sigma \in \Pi^h(s)\}.$$

Default gives yet another very important class of examples. The promises may be the same, but the deliveries may vary. In previous work [Dubey, Geanakoplos, and Shubik 2001] we supposed that agents incurred penalties per dollar of default. Since these penalties were taken to be idiosyncratic, and since agents differed in their marginal utilities of consumption and in their endowments, it turned out that they chose to deliver differently on the same promise.

It is helpful to consider one more example before turning to insurance. Let us imagine that farmers bring their heterogeneous fruit to the same bin. Buyers purchase by the pound, getting a uniform sample of the fruit in the bin. Even if each farmer knows the quality of his own fruit before he decides how much to sell, our analysis shows that one fruit price will clear the market. However, if the pool was violated, and buyers were allowed to reach in and pick out their favorite fruit, then equilibrium would break down. This is of more than academic interest. In the mortgage market, if the agencies, who are aware of the individual homeowner characteristics, were allowed to cherry pick the best loans, that gigantic market would also break down.

V. POOLING AND ADVERSE SELECTION WITH A SINGLE COOPERATIVE

To make our analysis concrete, we shall return frequently to the following canonical example and its straightforward generalization, which we shall call the microeconomic version of the insurance problem.

Let there be $H = 6$ household types, and $S = 3$ states of nature. Suppose that households have the same utility,

$$u^t(x_1, x_2, x_3) = \sum_{s=1}^3 \frac{1}{3} u(x_s), \text{ for all } t \in (0, 6],$$

where $u' > 0$, $u'' < 0$, and $\lim_{x \rightarrow 0} u'(x) = \infty$. The endowments of the households are given by

$$e^1 = \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}; e^2 = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}; e^3 = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix};$$

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

In our example, we always assume deliveries to be identical with endowments; i.e., $d^h = e^h$ for all $h \in \mathcal{H}$.

Households of type 1, 2, and 3 deliver two-thirds of the time, and are therefore called the *reliable class* R ; those of type 4, 5, and 6 deliver one-third of the time, and are called the *unreliable class* U .

Since households of each class are symmetric across states, we confine attention to equilibria in which $\varphi^t = \varphi^{t+1} = \varphi^{t+2}$, for all $t \in (0,1]$ and all $t \in (3,4]$. We do *not* assume that all reliable (or unreliable) households act the same, only that the triplets $(t, t+1, t+2)$ do. We wish to maintain the symmetry across triplets because when we interpret our example in the Rothschild-Stiglitz setting in Section IX, each of our triplets will correspond to a single Rothschild-Stiglitz household. Triplet symmetry implies that the pool has the same delivery rate in all states: $K_1 = K_2 = K_3 = \kappa$. The analysis collapses to a two-dimensional picture. Every household begins with an endowment of 1 in his "good" state(s) and 0 in his "bad" state(s). His final consumption x_s will only depend on whether s is good or bad for him. The reliable households $t \in (0,3]$ have utility of consumption

$$u^R(x_G, x_B) = \frac{2}{3} u(x_G) + \frac{1}{3} u(x_B),$$

while the unreliable households $t \in (3,6]$ have utility

$$u^U(x_G, x_B) = \frac{1}{3} u(x_G) + \frac{2}{3} u(x_B).$$

If the pool quality is $K = (\kappa, \kappa, \kappa)$, then by contributing θ , any household gives up θ in his good state, and receives $\theta\kappa$ in both states. His final consumption must therefore lie on the " θ -quantity line" starting at $(1 - \theta, 0)$ and moving northeast at 45° . (See Figure I.) On net he gives up $\theta(1 - \kappa)$ in his good state, and




FIGURE I

receives $\theta\kappa$ in his bad state. Thus, his consumption must also lie on the " κ -price line" joining $(1,0)$ to $(0,\kappa/(1-\kappa))$. The feasible consumption set, for a household contributing $0 \leq \theta \leq Q$ to a pool with quantity limit Q , is therefore the segment on the pool's κ -price line between $(1,0)$ and its intersection with the Q -quantity line.

From strict concavity of the utility functions, it is clear that each agent t has a unique optimal choice x^t on his feasible consumption set. This choice can be implemented by a unique θ^t . Thus, it is evident that if two households of the same class act exclusively on a common pool, they must act identically. Hence we can denote equilibrium with a single pool by $(\kappa, \varphi^U, \varphi^R) \in \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R}_+$.

If $Q \geq 1$, and the delivery rate $\kappa \in (0,1)$, then a household who maximizes $pu(x_G) + (1-p)u(x_B)$, with $p = \kappa$, will choose to contribute exactly one unit to the pool. In particular, if $\kappa = 1/3$ (or, $\kappa = 2/3$), then an unreliable (or a reliable) household will choose to contribute one unit and achieve his optimum on the unconstrained price line.

Reliable households are likely to curtail their contributions because they recognize that their promise delivers more on average than the pool, which is "debased" by the unreliable agents. When $\varphi^R < \varphi^U$, the pool delivery rate κ is worse than the




FIGURE II

population average of $1/2(1/3) + 1/2(2/3) = 1/2$, and we say that the pool displays adverse selection.

We can see pictorially why there is a tendency for adverse selection. If at some $\kappa > 0$ the reliable households voluntarily contribute $0 < \varphi^R < Q$, consuming $x^R = e^R + \varphi^R(K - e^R)$, then their indifference curve I^R through x^R must be tangent to the κ -price line. But the unreliable indifference curve I^U through x^R is flatter, and so the unreliable must be choosing $\varphi^U > \varphi^R$. Even if reliable households are up against the quantity constraint Q of the pool, it is evident from the single crossing property (depicted in Figure II) that the unreliable will find the constraint even more binding, so that, in any case, $\varphi^U \geq \varphi^R$. We thus have

LEMMA 1. Suppose that φ^U and φ^R are optimal contributions of unreliable and reliable households when they act exclusively on the same pool. Then $\varphi^U \geq \varphi^R$. Moreover, if $\varphi^R < Q$, then $\varphi^U > \varphi^R$.

In our canonical example with log utilities $u(x) = \log x$, and one pool with quantity constraint Q , it can easily be shown that there is a unique active equilibrium $(K(Q), \varphi^U(Q), \varphi^R(Q))$. The reader can verify that when Q is increased from 1 to $6/5$, every household is worse off. Although each unreliable household t wants to trade $\varphi^t = 6/5 > 1$, the upshot of all the unreliable

households doing so is to reduce the quality of the pool from $\kappa \approx .46$ to $\kappa \approx .44$ and to lower *every* household's utility, including their own.

Thus, the cooperative can help everybody by imposing a quantity restriction, $Q = 1$. Reducing Q further will help the reliable households and hurt the unreliable households. Reducing Q even further will hurt both households.

How will the cooperative set its quantity limit Q ?

VI. COMPETING COOPERATIVES

Different quantity limits may impinge on households differently. But if a cooperative cannot discriminate between households, it can set only one quantity limit. This gives an opportunity for a new cooperative to form, with a different quantity limit, to lure away dissatisfied members. How will this competition turn out?

Let us imagine a collection of cooperatives $j \in J = \{1, \dots, J\}$, all entailing (for simplicity) the same promises d_s^t , but different quantity restrictions $\phi_j^t \leq Q_j$. This defines the economy $((u^h, e^h, d^h)_{h \in \mathcal{H}}, (Q_j)_{j \in J})$.

Now household t chooses a vector $\theta = (\theta_1, \dots, \theta_J) \in \mathbb{R}_+^J$, where θ_j denotes the number of promises contributed to pool j . Suppose that for each pool j , the households anticipate deliveries K_{sj} in state s , per unit contributed. Denote $K_j = (K_{1j}, \dots, K_{Sj})$ and $K = (K_1, \dots, K_J)$. Household t then consumes

$$\chi^t(\theta, K) \equiv e^t + \sum_{j \in J} \theta_j (K_j - d^t).$$

He budget set is

$$\Sigma^t(K) = \{(\theta, y) \in \mathbb{R}_+^J \times \mathbb{R}_+^S : \theta_j \leq Q_j \text{ for all } j \in J, y = \chi^t(\theta, K)\}.$$

This is easily seen to be convex. But if we impose an exclusivity constraint as in Rothschild-Stiglitz, prohibiting any household from contributing to more than one pool, we obtain the nonconvex budget set,

$$\underline{\Sigma}^t(K) = \{(\theta, y) \in \Sigma^t(K) : \theta_j > 0 \Rightarrow \theta_k = 0 \text{ for all } k \in J \setminus \{j\}\}.$$

As we shall see, equilibrium—indeed refined equilibrium—exists in spite of this nonconvexity.

The notion of equilibrium can be extended in a straightforward manner to the setting of multiple, competing cooperatives.

Shorten "almost all t in $(0, H]$ " to "a.a.t.," and the integral $\int_0^H f(t) dt$ to \bar{f} . The vector $(K, \varphi, x) \in \mathbb{R}_+^{S \times J} \times \mathbb{R}_+^{J \times (0, H]} \times \mathbb{R}_+^{S \times (0, H]}$ is said to be an *equilibrium* if φ and x are measurable, and

$$(3) \quad K_{sj} = \frac{1}{\bar{\varphi}_j} \int_0^H \varphi_j^t d_j^t dt \text{ if } \bar{\varphi}_j > 0, \forall j \in J$$

$$(4) \quad (\varphi^t, x^t) \in \operatorname{argmax}_{(\theta, y) \in \Sigma^t(K)} u^t(y) \text{ for a.a.t.}$$

For simplicity, we have taken each agent's shares in the deliveries of a pool to be equal to his promises to the pool. We could have decoupled buying and selling by introducing a price π_j for shares of pool j , allowing an agent t to purchase shares $(\psi_1^t, \dots, \psi_J^t)$ across all pools, provided that $\sum_{j=1}^J \pi_j \psi_j^t = \sum_{j=1}^J \pi_j \varphi_j^t$. Since we will be focusing on the canonical insurance version of the model in which deliveries are uniform across states, i.e., $K_j = (\kappa_j, \dots, \kappa_j)$, this added flexibility actually adds nothing. In equilibrium we would have $\pi_j = \kappa_j$, and all agents would be content to choose $\psi^t = \varphi^t$ anyway (since at those prices, the shares of all pools are perfect substitutes). Thus, there is no loss of generality in dropping the π_j , ψ_j^t and supposing that sales determine purchases.

VII. COOPERATIVES WITHOUT MANAGERS, CONTRACTS WITHOUT DESIGNERS

In our framework the cooperative j makes no decisions. It simply stands open for business. Its quantity limit Q_j is its defining characteristic, rather than a strategic choice made by a manager. And its K_j is determined by the forces of perfect competition in equilibrium.

The current orthodox view is that insurance is impossible without strategic intermediaries, actively designing contracts. This view was most elegantly expressed by Rothschild and Stiglitz [1976], who described an economy with perfectly competitive consumers and oligopolistic, risk-neutral insurance companies. These companies designed and marketed insurance contracts (Q, π) specifying the quantity Q of insurance available and its price π . In their equilibrium, precisely two contracts, (Q, π) and $(\bar{Q}, \bar{\pi})$, are offered. To check its viability, every other potential contract (Q', π') is put on the market, one at a time, to see whether it can lure away a clientele with delivery rate $\kappa' > \pi'$.

Thus, customers never contemplate more than two or three contracts at the same time.

In our model, a household is presented with a full menu $\{(Q_j, \kappa_j = \pi_j)_{j \in J}\}$ where the set J can be arbitrarily large. The prices κ_j are highly nonlinear in Q_j . This sophistication is due entirely to "the market," not to any manager-designer. We will also see that only a few (Q_j, κ_j) have $\bar{\varphi}_j > 0$ among all potential $j \in J$. The set of active contracts, that are played out at equilibrium, is thus sharply determinate. And it is designed entirely by the "invisible hand" of perfect competition.

VIII. EQUILIBRIUM REFINEMENT

VIII.A. *The Definition*

With only one cooperative, we were content to confine our attention to equilibria in which the pool was active. With many cooperatives, the analogue would be to assume that all pools are active. But, as we have said, in the typical case every equilibrium effectively renders most pools inactive. Thus, we have no choice but to confront how anticipations K_j will be formed when pool j is inactive, since it is those anticipations themselves that are responsible for the inactivity.

Our definition permits any pool j to be inactive, i.e., to have $\bar{\varphi}_j = 0$. Many potential pools in the real world are also inactive. One possible explanation is that people anticipate unduly pessimistic deliveries from them and are thus discouraged from joining them. There is nothing so far in our definition to prevent this from happening. When pool j is active, there is a "reality check" on K_j , since (by (1)) K_j must conform to actual deliveries. But for inactive pools j , there are no real deliveries with which to compare K_j . If K_j were set suitably low, then no household t would be willing to contribute to pool j , for he would get very little per unit but incur a relatively large obligation to deliver d^t . Indeed, given an arbitrary subset of pools, we can always obtain equilibria which render them inactive by choosing their K_j to be low enough.

We believe that unreasonable pessimism does prevent many real world markets from opening, and provides an important role for government intervention. But it is interesting to study equilibrium in which anticipations 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.

Anticipated deliveries from inactive pools are analogous to beliefs in game theory "off the equilibrium path." Selten [1975] dealt with the game theory problem by forcing every agent to tremble and play all his strategies with probability at least $\epsilon > 0$, and then letting $\epsilon \rightarrow 0$. We shall also invoke a tremble, but in quite a different spirit. Our tremble will be "on the market" and not on households' (players') strategies. Indeed, no household could tremble the way we want: we introduce an external player who delivers more per unit than any of the real households.² This extraordinary delivery is what banishes whimsical pessimism.

Consider an external d -agent who contributes $\epsilon(n) = (\epsilon_j(n))_{j \in J} \geq 0$ to every pool, and delivers an exogenously fixed vector $d = (d, \dots, d)$ per unit contributed. We require that $d \geq \max_{h \in H} d_s^h$ for all $s \in S$. Any d satisfying this requirement will be called *optimistic*. The vector d indicates the boosting of household anticipations brought about by the external d -agent. We assume that $\epsilon(n) \rightarrow 0$ as $n \rightarrow \infty$, so one might interpret this agent as a government which guarantees delivery on the first infinitesimal promises.

The external d -agent may boost the delivery rate $K_{sj}(n)$ above the level achieved by the real households in the perturbed equilibrium $E_a(n)$. As $n \rightarrow \infty$, this boost disappears for pools that are active in the limit. But for inactive pools, his presence prevents the limiting anticipations from sinking too low, and steers them away from undue pessimism. In fact, at first glance, one might think that given his extraordinary deliveries, no pool will be inactive in equilibrium. We shall see, however, that quite the opposite is true: many pools will be inactive.

In the Appendix we explicitly add a d -external agent to the market (who contributes $\epsilon_j(n) > 0$ on every pool j and delivers $\epsilon_j(n)d$), and show that an $\epsilon(n) - d$ -equilibrium exists, and finally let $\epsilon(n) \rightarrow 0$ and take limits. This limit meets the criteria of the refined equilibrium we give below.

Computing a different equilibrium for each n , however, may be very time-consuming. Therefore, our definition of refinement captures the spirit of this limiting process, but makes the computation much easier by dropping $\epsilon_j(n) > 0$ unless $\bar{\varphi}_j(n) = 0$

2. If we were to invoke a tremble on strategies, e.g., forcing each household t to contribute $\epsilon > 0$ to every pool, this would *not* meet our needs.

(where the external boost $\epsilon_j(n) > 0$ is really needed) and also dropping the condition that $K_j(n) =$ actual deliveries for active j , since we know where these $K_j(n)$ must converge anyway. Our refinement is more permissive than that obtained by literally adding an external agent, but it is much easier to compute. Since we shall prove uniqueness of our refined equilibrium, its expansive definition may also be taken to be an advantage.

Formally, we say that an equilibrium $E = (K, \varphi, x) \in \mathbb{R}_+^{S^J} \times \mathbb{R}_+^{J \times (0, H)} \times \mathbb{R}_+^{S \times (0, H)}$ is a *refined equilibrium* if there is a sequence $E_q(n) = (K(n), \varphi(n), x(n), \epsilon(n)) \in \mathbb{R}_+^{S^J} \times \mathbb{R}_+^{J \times (0, H)} \times \mathbb{R}_+^{S \times (0, H)} \times \mathbb{R}_+^J$ such that d is optimistic, $\varphi(n)$ and $x(n)$ are measurable for all $n = 1, 2, \dots$ and

$$(5) \quad \epsilon(n) \rightarrow 0, K(n) \rightarrow K \text{ and } \varphi^t(n) \rightarrow \varphi^t, x^t(n) \rightarrow x^t \text{ for a.a. } t$$

$$(6) \quad (\varphi^t(n), x^t(n)) \in \operatorname{argmax}_{(\theta, y) \in \Sigma^t(K(n))} u^t(y) \text{ for a.a. } t \text{ and all } n$$

$$(7) \quad \epsilon_j(n) > 0 \text{ if } \bar{\varphi}_j(n) = 0, \text{ for all } j \in J \text{ and all } n,$$

$$(8) \quad \text{For all } n \text{ and for all } j \in J^* = \{j \in J : \bar{\varphi}_j = 0\},$$

$$K_{sj}(n) = \frac{1}{\epsilon_j(n) + \bar{\varphi}_j(n)} \left[\epsilon_j(n) d + \int_0^H \varphi_j^t(n) d_s^t dt \right].$$

THEOREM 1. Consider the finite type continuum model with competing cooperatives and the exclusivity constraint. Then a refined equilibrium always exists.

Proof of Theorem 1. See the Appendix.

Theorem 1 shows that adverse selection and signaling do not compromise the universal existence of perfectly competitive equilibrium, putting to rest whatever doubts might still linger on the subject. In Section IX we recast the Rothschild-Stiglitz insurance model as a special case, thus showing that it too must always have an equilibrium. In Sections X and XI we explicitly compute the insurance equilibrium and show that it is unique. We find that it retains the economic flavor of the Rothschild-Stiglitz separating equilibrium.

VIII.B. Expectations at Inactive Pools in Refined Equilibrium

Consider again our canonical example, with reliable households who always deliver 2/3 and unreliable households who always deliver 1/3, but with J pools $(Q)_{j \in J}$. Suppose that we have




FIGURE III

an equilibrium $(\kappa, \varphi^R, \varphi^U)$, where $\kappa = (\kappa_j)_{j \in J}$ and $K_j = (\kappa_j, \dots, \kappa_j)$, in which pool J is inactive. We inquire what values κ_J could take to “sustain” pool J , i.e., keep it inactive while satisfying the refinement condition.

Holding $(\kappa_j)_{j=1}^{J-1}$ fixed, suppose that the supply curve of contributions to pool J is given by Figure III, where $0 < 1/3 \leq \underline{\kappa} < \bar{\kappa} \leq 2/3$. Suppose that unreliable households start contributing when $\kappa_J = \underline{\kappa}$, while reliable households only start at $\kappa_J = \bar{\kappa}$. Suppose furthermore that, at $\kappa_J = \underline{\kappa}$, unreliable households are indifferent between contributing any quantity from 0 to $q > 0$.

Without the equilibrium refinement, any $\kappa_J < \underline{\kappa}$ is sustaining. But our equilibrium refinement requires that if $\kappa_J < \underline{\kappa}$, then $\kappa_J \geq 1$, since no perturbation would induce contributions apart from the external agent, who delivers $d \geq \max_{h \in H, s \in S} d_s^h = 1$. This contradiction shows that if κ_J is sustaining, then $\kappa_J = \underline{\kappa}$. (Clearly, $\kappa_J > \underline{\kappa}$ is not sustaining, since the unreliable would be selling.)

Indeed, $\kappa_J = \underline{\kappa}$ is sustaining. At that level, everybody is willing not to contribute. Moreover, at any small perturbation, only the unreliable can be induced to contribute to J , and their delivery rate is $1/3 \leq \kappa_J$, so κ_J is not unduly pessimistic.

To formally check that $\kappa_J = \underline{\kappa}$ satisfies the perturbation requirement, shift a small population μ of U households to pool J , together with a small contribution $\epsilon(\mu)$ from the external agent, satisfying

$$\frac{\mu(1/3) + \epsilon(\mu) d}{\mu + \epsilon(\mu)} = \kappa,$$

and let the rest of the population act as in the equilibrium. Keep the delivery rates $(\kappa_j(\mu))_{j \in J} = (\kappa_j)_{j \in J}$ for all μ . Since $1/3 \leq \kappa < d$, there exists a unique solution $\epsilon(\mu) \geq 0$; and, as $\mu \rightarrow 0$, $\epsilon(\mu) \rightarrow 0$. Note that every household is still optimizing in the perturbation and all but the small measure μ are acting exactly as in the equilibrium, so all the conditions of refinement are met.

The fact that it is possible to sustain no trade on pool J in refined equilibrium is the reason why existence always prevails in our model.

Note that such a simple perturbation worked because the supply curve for unreliable agents is flat at κ . This allowed us to keep all the expectations $(\kappa_j(\mu))_{j \in J}$ fixed at $(\kappa_j)_{j \in J}$ and still produce contributions on pool J . If the supply curve of the unreliable agents had not been flat, we would have needed a more delicate perturbation, with $\kappa_j(\mu) \neq \kappa_j$. Furthermore, if we had defined refined equilibrium in terms of an external agent who trembles positively on every market, we would again need another perturbation, with $\kappa_j(\mu) \neq \kappa_j$, when $\kappa = 1/3$, since in that case our perturbation defined $\epsilon(\mu) = 0$, which would no longer be admissible. Thus, the simplicity of our perturbation also stems from the fact that the external agent is not required to act on a pool if real households are already acting there.

There are circumstances in which the refinement rules out inactivity on pool J . That is the reason we are able to prove uniqueness of equilibrium in the insurance examples to come. For instance, consider the situation where the unreliable and reliable supply curves are reversed, so that it is the reliable households who begin contributing at the lower price κ , and unreliable households who begin to contribute at the higher rate $\bar{\kappa} \leq 2/3$. As before, the only expectation that could sustain inactivity is $\kappa_J = \kappa$. But any perturbation around κ only induces trade of reliable households, plus perhaps the external agent. Hence deliveries must be at least $2/3$. But $\kappa < 2/3$, a contradiction.

VIII.C. Rothschild-Stiglitz Expectations

Consider again the situation depicted in Figure III, for which we found $\kappa_J = \kappa$ sustaining. Rothschild and Stiglitz would disagree. If expectations on pool J are dramatically raised to $\kappa^* > \bar{\kappa} > \kappa$, then reliable agents will want to contribute to pool J as

well. Adding their deliveries to the unreliable deliveries could give average deliveries which might indeed be as high as κ^* , and if so, Rothschild and Stiglitz argue that pool J cannot be inactive. Or else, if pool J is inactive, they suggest another pool $J + 1$ would be created with $Q_{J+1} = Q_J$ and $\kappa_{J+1} = \kappa^*$ that would break the equilibrium. In short, their view is that every contract (Q, κ) must be reckoned with, and (Q_J, κ^*) destroys the existence of equilibrium.

Our model works differently. The contract or pool terms Q_j are set exogenously, and might be all inclusive (take J to include every quantity level). However, *one* expectation $\kappa_j(Q_j)$ is then formed for each pool so that all the markets clear. No household need ever consider an expectation $\kappa'_j \neq \kappa_j(Q_j)$. Our perfectly competitive model therefore has far fewer contracts than the Rothschild-Stiglitz model. But it still has enough contracts to incorporate the economics of adverse selection and signaling. Indeed, we shall prove that in the Rothschild-Stiglitz insurance setting, refined equilibrium is unique.

Rothschild and Stiglitz might have argued that instead of thinking of the pools as strategic dummies, we could imagine that they were each run by some entrepreneur. He might personally guarantee deliveries of κ^* , thereby forcing households to consider this expectation, and disrupting the old equilibrium.

This seems implausible to us in a large economy. We suppose that households are aware of the composition of deliveries at the prevailing equilibrium, and perhaps of how the composition would actually change if delivery expectations were slightly perturbed. But households lack the knowledge or computing power to infer the composition when expectations are far from the equilibrium. So we are led to wonder, how confident can the entrepreneur really be that the composition of willing contributors would be so radically better on account of his guarantee? We have in mind a competitive world with many small agents. If the little entrepreneur's gambit is to be successful, he must lure new reliable households at κ^* , who were unwilling to contribute at κ_J . But it is the unreliable, already willing to contribute at κ_J , who will be even more eager to contribute at κ^* and likely to get to him first. If so, his meager wealth will certainly not be enough to stand guarantee for his exorbitant offer of κ^* , and he will suffer a disaster. A cautious entrepreneur would forecast that the mix of deliveries elicited by him at κ^* is not going to be much better than what prevails in equilibrium at κ_J , and so would not undertake

the disastrous κ^* gambit. But even if he abandons caution, whether he is right or not is not in his hands. Why should a reliable household, prudently cautious unlike the entrepreneur, not forecast the disaster for herself and abstain from going to his pool (Q_{J,κ^*})? The pool offers the same terms Q_J as the equilibrium pool (Q_{J,κ_J}), plus a meaningless guarantee of κ^* without the wealth to back it. After all, if the guarantee breaks down and deliveries turn out to be not much better than κ_J , reliable people like herself stand to lose since they were not willing to contribute to (Q_{J,κ_J}) in the first place. Thus, the contributors most at risk from entering the pool (Q_{J,κ^*}), and therefore most unlikely to come to it, are precisely the ones the entrepreneur must count on to improve the deliveries and to help sustain his offer of κ^* . The entrepreneur should think again. If he has concern that there may be caution among his clientele, this should cause him to become cautious himself.

If the entrepreneur had enormous wealth, then his guarantee would be meaningful, and he might well induce all the reliable households to contribute, relieving himself of the need to use his wealth, and disrupting the κ_J equilibrium. But such large wealth is a feature of an oligopolistic model. We rigorously maintain the hypothesis of perfect competition.

VIII.D. *The Optimistic External Delivery*

We have chosen $d \geq \max_{h \in H} \max_{s \in S} d_s^h \equiv M$ because it keeps active as many pools as possible, and thus eliminates as many equilibria as possible. We show in the Appendix, subsection B, that, in the context of our insurance model, the set of refined equilibria $\mathcal{E}(d)$ is inversely monotonic: $d > d'$ implies that $\mathcal{E}(d) \subseteq \mathcal{E}(d')$. But $\mathcal{E}(d)$ is essentially the same as $\mathcal{E}(d')$ if $d > d' \geq M$ (only the equilibrium delivery rates for inactive pools might be affected). Thus, we have chosen our perturbation to make existence as hard as possible, and uniqueness as easy as possible.

IX. INSURANCE

The classical insurance problem can be embedded in our model of cooperatives, and turns out to be a straightforward generalization of our canonical example.

IX.A. The Rothschild-Stiglitz Insurance Problem

As in Rothschild-Stiglitz, we consider a continuum of two types of households: "reliable" (R) and "unreliable" (U), with population measures λ_R and λ_U , respectively. Each household knows his own type, but not that of the others. Each household 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 households. The unreliable households are more accident-prone than the reliable. Thus, if p^h denotes the probability of a good state for type h , we have $p^R > p^U$.

The assumption of a continuum of households with independent accidents is actually quite strong. It implies that the same proportion of *any* nonnull subset of households of a given class has an accident in almost every state.

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

$$p^h u(x_G) + (1 - p^h) u(x_B),$$

to a household of type $h = R, U$. We take p^h to be a rational number m/n .

IX.B. The Canonical Insurance Model

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 that make explicit who has an accident there. This makes it clear that "identical" insurance policies for two households of the same class do not pay off identically, since the households will have accidents in different states, even if their probabilities are the same.

Within our framework of finite states and household types, 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 $1 - p^R = r/n$, and let $1 - p^U = u/n$. To embed the insurance problem in our framework,

take $S = n$. Let the intervals $(0, \lambda_R]$ and $(\lambda_R, \lambda_R + \lambda_U]$ represent the reliable and unreliable households. Partition the reliable households into $\binom{n}{r} = n!/r!(n-r)!$ consecutive intervals, each of length $\alpha_R \equiv \lambda_R/\binom{n}{r}$, and similarly divide the unreliable households into $\binom{n}{u}$ intervals, each of length $\alpha_U \equiv \lambda_U/\binom{n}{u}$. These intervals correspond to types in our canonical model. Each type τ can be identified with a distinct subset $S_\tau \subset S$ of bad states (r in number if reliable, u in number if unreliable). All households t of type τ have endowments equal to 1 if $s \in S \setminus S_\tau$, and equal to 0 if $s \in S_\tau$.

The reader can verify that each household has the right probability of accident (r/n if reliable, u/n if unreliable), and that in every state $(r/n)\lambda_R$ of reliable and $(u/n)\lambda_U$ of unreliable households have accidents.

Rothschild and Stiglitz have implicitly assumed much more, namely that for any nonnull set of reliable agents, precisely r/n of them will have an accident in almost every state (and similarly u/n for unreliable). We reproduce this feature of their model by identifying each of their reliable households $s \in (0, \lambda_R]$ with a $\binom{n}{r}$ -tuple $(t, t + \alpha_R, t + 2\alpha_R, \dots, t + [\binom{n}{r} - 1]\alpha_R)$ for $t = s/\binom{n}{r} \in (0, \alpha_R)$. Similarly, each Rothschild-Stiglitz unreliable household $s + \lambda_R \in (\lambda_R, \lambda_R + \lambda_U]$ corresponds to a $\binom{n}{u}$ -tuple $(t, t + \alpha_U, t + 2\alpha_U, \dots, t + [\binom{n}{u} - 1]\alpha_U)$ for $t = (s/\binom{n}{u}) + \lambda_R \in (\lambda_R, \lambda_R + \alpha_U]$. The reader can verify that for any nonnull subset of reliable Rothschild-Stiglitz households, precisely the fraction r/n of the corresponding reliable households in our model have an accident in every state $s \in S$ (and similarly for the unreliable).

To keep the identity of the Rothschild-Stiglitz household inviolate, we confine attention to the scenario where all the members of any given tuple behave identically, i.e., $\varphi^t = \varphi^{t+\alpha_R}$ if t and $t + \alpha_R$ are reliable (and similarly $\varphi^t = \varphi^{t+\alpha_U}$ if t and $t + \alpha_U$ are unreliable). This generalizes the triplet-symmetry of our example. We shall call it "tuple-symmetry," and assume that it holds not just in equilibrium but also in deviation from equilibrium. (Since any tuple has zero measure, a deviation by it has no effect on the continuum.) Consider now our equilibrium refinement. From tuple-symmetry and the fact that the external agent delivers identically across states, it follows that $K_j(\epsilon)$ has identical components across states in any ϵ -perturbation of equilibrium. Therefore, $K_j = \lim_{\epsilon \rightarrow 0} K_j(\epsilon)$ also inherits this property in any refined equilibrium (obtained as a limit of the perturbations). This fact will be assumed from now on in our analysis of the

canonical insurance model: for any pool j , we shall always take $K_j \equiv (\kappa_j, \dots, \kappa_j)$ to have the same component κ_j in every state $s \in S$.

Recalling our numerical example of Section V, note that it corresponds to the insurance problem with $p^R = 2/3$, $p^U = 1/3$, $\lambda_R = \lambda_U = 3$. Hence, in the microeconomic representation provided by our example, $S = 3$. There are $\binom{3}{1} = 3$ reliable types, each of measure $3/\binom{3}{1} = 1$, and $\binom{3}{2} = 3$ unreliable types, each of measure $3/\binom{3}{2} = 1$.

X. EXISTENCE OF THE SEPARATING EQUILIBRIUM IN THE CANONICAL INSURANCE MODEL

Rothschild and Stiglitz [1976] made the important observation that adverse selection in insurance markets might lead to the same kind of inefficient signaling 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, some 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 signaling there might not be any equilibrium in insurance markets.

The only equilibrium Rothschild and Stiglitz found is the "separating" equilibrium in which each household class takes out a different insurance. Rothschild and Stiglitz observed that in such an equilibrium the unreliable class should feel unconstrained by the quantity restriction while the reliable class should feel quantity constrained. Moreover, the unreliable types should be indifferent to either of the two contracts, while the reliable class should strictly prefer their quantity constrained contract.

We get the same sort of equilibrium, though not quite exactly because our menu of quantity constraints is finite. (They examined the simpler case where all quantity constraints are available.) The important difference is that our separating equilibrium always exists, whereas Rothschild-Stiglitz found robust regions of nonexistence.

Theorem 1 already guarantees that some equilibrium always exists in our general pooling model. Specializing to the canonical




FIGURE IV

insurance model, we are able to go a step further, and precisely describe what equilibrium must look like. To this end we give a constructive proof of existence (and uniqueness) which does not rely on Theorem 1.

We focus on the example described in Section V. This is for concreteness. The arguments can easily be modified to give the same results for the canonical insurance model.

We shall make heavy use of Figure I in what follows. For any consumption $x = (x_G, x_B)$ with $x_G \leq 1$, define $\varphi(x)$ to be the quantity line passing through x (so $\varphi(x) = 1 - x_G + x_B$); let $\kappa(x)$ be the price line passing through x (so $\kappa(x) = x_B/\varphi(x)$). Let $I^R(x)$ and $I^U(x)$ be the reliable and unreliable indifference curves through x .

Let Z^U be the optimal consumption of U -households on the (unconstrained) 1/3-price line, and $I^U = I^U(Z^U)$ be their indifference curve through Z^U . (See Figure IV.) Since $p^U = 1/3$, $\varphi(Z^U) = 1$. Let Z denote the intersection of I^U with the 2/3-price line, and $\varphi(Z)$ the quantity line through Z . Let $I^R = I^R(Z)$ be the R -indifference curve through Z .

Let $\mathbb{Q} = \{Q_1, \dots, Q_J\}$ with $0 \leq Q_1 \leq \dots \leq Q_J$ and $Q_J \geq 1$, denote the grid of quantity signals.

THEOREM 2.1 (Existence of Fully Separating Equilibrium). If $\varphi(Z) = Q_{j^*} \in \mathbb{Q}$, then the canonical insurance model has a refined equilibrium in which all reliable households contribute $\varphi^R = \varphi(Z)$ to pool $j^* < J$, and all unreliable households contribute $\varphi^U = 1$ to pool J .

Observe that the existence of a fully separating equilibrium requires only that \mathbb{Q} contain at least two signals, $\varphi(Z)$ and $Q_j \geq 1$. So, if we (like Rothschild-Stiglitz) had simply taken \mathbb{Q} to be the comprehensive set of all signals, i.e., $\mathbb{Q} = [0, \infty)$, then Theorem 2.1 would always apply. We have opted for a finite grid \mathbb{Q} because it makes for a notationally simpler analysis. It also raises the interesting question of what would happen if $\varphi(Z) \notin \mathbb{Q}$.

When $\varphi(Z) \notin \mathbb{Q}$, we need to be sure that there are signals $Q_k < \varphi(Z) < Q_{k+1}$ very close to $\varphi(Z)$. This will necessarily be true if the *grid size* $\delta(\mathbb{Q})$ of the quantity signals \mathbb{Q} ,

$$\delta(\mathbb{Q}) \equiv \max \{Q_1, \max_{1 < j < J} (Q_{j+1} - Q_j)\}$$

is small; i.e., the grid is dense.

The following theorem holds even if $\varphi(Z) \notin \mathbb{Q}$.

THEOREM 2.2 (Existence of Nearly Separating Equilibrium).

There exist $\delta > 0$ and $c > 0$ such that if $\delta(\mathbb{Q}) \leq \delta$, then there is a refined equilibrium in which all reliable households contribute Q_{j^*} to a pool $j^* < J$ with $|\varphi(Z) - Q_{j^*}| < \delta(\mathbb{Q})$, and at least the fraction $1 - c\delta(\mathbb{Q})$ of all unreliable households contribute 1 to pool J . The remaining unreliable households, if any, contribute Q_{j^*} on j^* , creating at most a small degree of heterogeneous pooling on j^* .

Even in the missing quantity case where $\varphi(Z) \notin \mathbb{Q}$, a purely separating equilibrium often exists. But sometimes it becomes necessary to split the unreliable households between two pools. In the proof of Theorem 2.2 we pinpoint conditions under which splitting must occur. Theorem 2.2 shows, however, that the splitting goes to zero linearly with the grid size $\delta(\mathbb{Q})$. The splitting has nothing to do with the proportion of reliable and unreliable households in the whole population, and so nothing to do with the nonexistence of Rothschild-Stiglitz equilibrium (which depended on there being a high proportion of reliable households).

X.A. Proof of Theorem 2.1 (The Case $\varphi(Z) \in \mathbb{Q}$)

We shall first prove existence of a separating equilibrium under the assumption that there is a pool j^* such that $\varphi(Z) = Q_{j^*}$. And, in our discussion of refinement, we shall suppose that the d -external agent is chosen with $d = 1$.

Let the unreliable households contribute $\varphi_J^U = 1$ to pool J , with $\kappa_J = 1/3$, to obtain their optimal consumption Z^U on the




FIGURE V

1/3-price line. (This is feasible since $Q_j \geq 1$.) Let the reliable households contribute $\varphi_{j^*}^R = Q_{j^*}$ units to pool j^* , with $\kappa_{j^*} = 2/3$, to obtain the consumption $Z^R = Z$. Note that $Q_{j^*} < 1$ and that reliable households would contribute one unit on the unconstrained 2/3-price line. So Q_{j^*} is indeed the optimal contribution for them on pool j^* . Furthermore, it follows that the I^R curve cuts the I^U curve transversally from "above-left" at Z^R , as shown in Figure V (staying above I^U on the left of Z^R , and below it on the right).

We still must price all the inactive pools in a manner that satisfies our equilibrium refinement.

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

If $Q_{j^*} < Q_j < 1$, then the Q_j -quantity line intersects I^U at Z^U , before it intersects I^R . Set $\kappa_j = \kappa(Z^U)$ in accordance with the dotted line in Figure V, which connects (1,0) to Z^U .

If $Q_j < Q_{j^*}$, then the Q_j -quantity line intersects I^R at $Z^{\tilde{R}}$ before it intersects I^U . If $Z_G^R < 1$, set $\kappa_j = \kappa(Z^{\tilde{R}})$ in accordance with the other dotted line in Figure V, which connects (1,0) to $Z^{\tilde{R}}$.

If $Q_j < Q_{j^*}$ and $Z_G^R \geq 1$, set $\kappa_j = 1$.

For every $j \notin \{J, j^*\}$ we show that no household can improve by using pool j .

When $1 < Q_j$ (and so $\kappa_j = 1/3$), the unreliable households are




FIGURE VI

indifferent to contributing (one unit to) pool j and (one unit to) pool J , while the reliable households are strictly worse off using j .

When $Q_{j^*} < Q_j \leq 1$, the unreliable households are indifferent to contributing (Q_j units to) pool j and (one unit to) pool J , while reliable households are strictly worse off using j .

When $Q_j < Q_{j^*}$ and $\kappa_j < 1$, the reliable households are indifferent to contributing (Q_j units to) pool j and (Q_{j^*} units to) pool j^* , while the unreliable households are strictly worse off using pool j .

Finally, when $Q_j < Q_{j^*}$ and $\tilde{Z}_G^R \geq 1$, denote by \tilde{Z} the intersection of the 1-price line with the Q_j -quantity line. (See Figure VI.) Since I^R is downward sloping, \tilde{Z} is on or below I^R . As for the unreliable households, they are strictly better off at Z^U .

This verifies that we have defined an equilibrium.

We shall now construct the perturbation $E(n) = (\kappa(n), \varphi(n), x(n), \epsilon(n))$ which shows that our equilibrium is, in fact, refined.

Let $\kappa_j(n) = \kappa_j$ for all $j \in J$ and $n = 1, 2, \dots$

We next define $\varphi(n)$ by switching disjoint sets of households from their equilibrium actions onto inactive pools. Then we check that the deliveries in the perturbation are equal to $\kappa_j(n)$, for pools j that were inactive in the equilibrium. (Recall that the refine-

ment does *not* require that $\kappa_j(n)$ conform to actual deliveries for pools j that were already active in equilibrium. But in the perturbation we are going to describe, they will in fact conform.)

Take j with $Q_j \geq 1$. Let a (needless to say, tuple-symmetric) set of U -households of measure $1/n$ switch out of pool J and contribute one unit on pool j instead; and put $\epsilon_j(n) = 0$. Note that these households still consume their optimal Z^U , and their deliveries on pool j justify $\kappa_j(n) = 1/3$.

Next take j with $Q_{j^*} < Q_j < 1$. Let a set of U -households of measure $1/n$ switch out of pool J and contribute Q_j on pool j . Since $1/3 < \kappa_j < 2/3$, we can choose $\epsilon_j(n) > 0$ to satisfy

$$\frac{(1/n)Q_j(1/3) + \epsilon_j(n)}{(1/n)Q_j + \epsilon_j(n)} = \kappa_j = \kappa_j(n).$$

Again household optimality holds, and the $\kappa_j(n)$ stands justified.

Then take j with $Q_j < Q_{j^*}$. First suppose that $Z_G^R < 1$. Let a set of R -households of measure $1/n$ switch out of pool j^* and contribute Q_j on pool j . Since $2/3 < \kappa_j < 1$, we can choose $\epsilon_j(n) > 0$ to satisfy

$$\frac{(1/n)Q_j(2/3) + \epsilon_j(n)}{(1/n)Q_j + \epsilon_j(n)} = \kappa_j = \kappa_j(n).$$

Once again household optimality holds, and $\kappa_j(n)$ is justified.

Finally if $Q_j < Q_{j^*}$ and $Z_G^R \geq 1$, let $\epsilon_j(n) = 1/n$. Only the external agent acts on pool j in this case, and since he delivers $d = 1$, $\kappa_j(n) = 1$ is justified.

The definition of $x(n)$ follows from $\varphi(n)$ and $\kappa(n)$.

It is evident that our perturbation validates E as a refined equilibrium. QED

X.B. Proof of Theorem 2.2 (The Case $\varphi(Z) \notin Q$)

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

Let Q_k, Q_{k+1} be consecutive quantities in the grid Q which trap the missing quantity $\varphi(Z)$ in between. Denote

$W \equiv$ intersection of I^U with the Q_{k+1} -quantity line

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

We have already seen in Figure IV that the $I^R(Z)$ curve cuts the $\kappa(Z) = 2/3$ -price line transversally at Z from bottom right to top left. By continuity, this must remain true of the $I^R(W)$ curve




FIGURE VII

and the $\kappa(W)$ -price line for all W sufficiently close to Z , and hence for sufficiently small grid size, $\delta(Q) < \delta$, for some $\delta > 0$.

CASE 1 (Figure VIII). The reliable households (weakly) prefer Y to W .

Then define $Z^R = Y$ and $j^* = k$, and proceed exactly as before to price the inactive pools and to construct the perturbation. (Note that $I^R = I^R(Y)$ is above $I^R(W)$, which in turn stays strictly above the $\kappa(W)$ -line from $(1,0)$ to W , since $\delta(Q) < \delta$, and our arguments from before are not impeded.)

CASE 2 (Nearly Separating Equilibrium). (Figure IX) The reliable prefer W to Y .

In this case we do not get a pure separating equilibrium, but an equilibrium with a little splitting. Let $j^* = k + 1$. Let all the reliable households contribute $Q_{k+1} \equiv Q_{j^*}$ units to pool $k + 1 \equiv j^*$ and consume $W \equiv Z^R$. Set $\kappa_{j^*} = \kappa(W)$, in accordance with the dotted line joining $(1,0)$ to W in Figure IX. Recall that the dotted line stays below $I^R(W)$ when the grid $\delta(Q) < \delta$.

The new feature of this equilibrium is that some unreliable households also contribute to pool j^* . In fact, just enough of them contribute to j^* so that the delivery rate falls from $2/3$ to κ_{j^*} . The




FIGURE VIII

rest of the U population acts as before, contributing $\phi_J^U = 1$ units to pool J . The pricing of inactive pools and the perturbation work exactly as before.

The reader can check that the $\kappa(W)$ converges to $2/3$ linearly




FIGURE IX

with $\delta(\mathbb{Q}) \rightarrow 0$. Hence the splitting (i.e., measure of unreliable households contributing to j^*) also converges to zero linearly with $\delta(\mathbb{Q})$. QED

XI. UNIQUENESS OF EQUILIBRIUM

We prove that the separating equilibrium (or nearly separating equilibrium when $\varphi(Z) \notin \mathbb{Q}$) is unique. Rothschild and Stiglitz did not need to worry about the possibility that different households in the same class might split up and contribute to different pools. We allow for such possibilities, yet we will manage to prove uniqueness.

THEOREM 3.1 (Uniqueness of Fully Separating Equilibrium). If $\varphi(Z) \in \mathbb{Q}$, then there is a $\bar{\delta} > 0$ such that when $\delta(\mathbb{Q}) \leq \bar{\delta}$, any two refined equilibria (κ, φ, x) and $(\tilde{\kappa}, \tilde{\varphi}, \tilde{x})$ are essentially the same, satisfying

- (a) $\kappa_j = \tilde{\kappa}_j$ if $\kappa_j < 1$
- (b) $x^t = \tilde{x}^t$ for a.a. t
- (c) $\varphi^t = \tilde{\varphi}^t$ for a.a.(reliable) $t \in (0, 3]$
- (d) $\varphi_j^t > 0 \Rightarrow \varphi_j^t = 1$ and $\kappa_j = 1/3$ for a.a.(unreliable) $t \in (3, 6]$.

The only touch of nonuniqueness is irrelevant. In (d) unreliable households may choose pools j which effectively differ in name only. In (a), *inactive* pools with very low quantity limits may have different delivery rates $\kappa_j \geq 1$, stemming from the fact that different external d -agents have been involved in the perturbation used for refinement. (If we were to *fix* d , all these different delivery rates would be d , and (a) would hold for all $j \in J$.) It is clear that the ambiguities "at the edges" (large Q_j and small Q_j) have no real effect on the equilibrium.

When $\varphi(Z) \notin \mathbb{Q}$, the nearly separating equilibrium computed in Section X is also essentially unique. Recall Figure VII, in which W is defined as the intersection of I^U and the smallest quantity line above $\varphi(Z)$, and Y is defined as the intersection of the 2/3-price line with the largest quantity line below $\varphi(Z)$. We have

THEOREM 3.2 (Uniqueness of Nearly Separating Equilibrium).

There exists $\bar{\delta} > 0$ such that when $\delta(\mathbb{Q}) \leq \bar{\delta}$, any two refined equilibria (κ, φ, x) and $(\tilde{\kappa}, \tilde{\varphi}, \tilde{x})$ are essentially the same. If




FIGURE X

$u^R(Y) > u^R(W)$, then the conclusion of Theorem 3.1 is valid.
If $u^R(W) > u^R(Y)$, then

- (a) $\kappa_j = \tilde{\kappa}_j$ if $\kappa_j < 1$;
- (b) $x^t = \tilde{x}^t$ and $\varphi^t = \tilde{\varphi}^t$ for a.a.(reliable) $t \in (0, 3]$;
- (c) $\lambda(\{t \in (3, 6] : (\varphi^t, x^t) = (\theta, y)\}) = \lambda(\{t \in (3, 6] : (\tilde{\varphi}^t, \tilde{x}^t) = (\theta, y)\})$ for all $(\theta, y) \in \mathbb{R}_+^J \times \mathbb{R}_+^2$, where λ denotes Lebesgue (population) measure (In other words, the distribution of action-consumption pairs is the same.);
- (d) $u^t(x^t) = u^t(\tilde{x}^t)$ for a.a. t (obvious from (b) and (c)).

If $u^R(W) = u^R(Y)$, then (a) and (d) still hold.

XI.A. Proof of Theorem 3.1 (The Case $\varphi(Z) \in Q$)

Reconsider Figure IV with the 1/2-price line superimposed. Let Ω be the shaded region between I^R and the 1/2-price line. See Figure X. (Ω may be empty.)

Note that the region Ω is at a positive distance from the 2/3-price line. So there is a $\bar{\delta} > 0$ such that whenever $\delta(Q) \leq \bar{\delta}$, we have

- (*) For any $x \in \Omega$, there exists a pool j with $Q_j < \varphi(x)$ such that the Q_j -quantity line intersects $I^R(x)$ before it intersects the 2/3-price line. (See Figure XI.)

We present the proof of Theorem 3.1 through a sequence of




FIGURE XI

lemmas. Throughout, an arbitrary equilibrium (κ, φ, x) is assumed fixed, and the lemmas describe its various features (using the definitions of I^R, I^U, Z^U, Z , and $\varphi(Z) = Q_{j^*}$ from Figure IV).

LEMMA 2. All U -households obtain at least the utility given by I^U .

Proof of Lemma 2. By our refinement $\kappa_j \geq 1/3$. But I^U is the utility they could guarantee via pool J (with its generous quantity constraint $Q_j \geq 1$) if its delivery rate were just $1/3$. QED

LEMMA 3. $\kappa_j \geq 2/3$ if $j \leq j^*$.

Proof of Lemma 3. If $\kappa_j < 2/3$, then the U -households obtain strictly less utility than I^U via pool j . Hence (by Lemma 2) they abstain from pool j in any (small enough) perturbation used in the equilibrium refinement. But then, by the definition of our refinement, $\kappa_j \geq 2/3$. QED

LEMMA 4. All reliable households obtain at least the utility given by I^R .

Proof of Lemma 4. They can obtain consumption at least as good as Z via pool j^* , in view of Lemma 3. QED

Let λ_j^U, λ_j^R be the measure of the sets of U, R households who are active on (i.e., contribute positively to) pool $j \in J$.




FIGURE XII

LEMMA 5 (Restriction on Heterogeneous Pooling). There does not exist any pool j such that $\lambda_j^U \geq \lambda_j^R > 0$.

Proof of Lemma 5. Suppose that there is such a pool j . Then, by Lemma 1, $\varphi_j^U \geq \varphi_j^R$. So

$$\kappa_j = \left(\lambda_j^U \varphi_j^U \frac{1}{3} + \lambda_j^R \varphi_j^R \frac{2}{3} \right) / (\lambda_j^U \varphi_j^U + \lambda_j^R \varphi_j^R) \leq \frac{1}{2} \cdot \frac{1}{3} + \frac{1}{2} \cdot \frac{2}{3} = \frac{1}{2}.$$

Let \bar{Z}^R, \bar{Z}^U be the consumption of reliable, unreliable households contributing to pool j . By Lemma 4 and $\kappa_j \leq 1/2$, $\bar{Z}^R \in \Omega$. Let Q_i be the quantity just below $\varphi(\bar{Z}^R)$. From (*), we know that the Q_i -quantity line cuts $I^R(\bar{Z}^R)$ at a point Y before it cuts the $2/3$ -price line. From the single crossing property, and the fact $\varphi(\bar{Z}^U) \geq \varphi(\bar{Z}^R)$, we know that Y lies below $I^U(\bar{Z}^U)$. (See Figure XII.)

Let $\bar{\kappa}_i = \kappa(Y)$ correspond to the slope of the line (shown dotted in Figure XII) that joins $(1,0)$ to Y . By (*), $\bar{\kappa}_i < 2/3$. We shall show that it is impossible to assign an equilibrium delivery rate κ_i to pool i . If $\kappa_i \leq \bar{\kappa}_i$, then the κ_i -price line and the Q_i -quantity line intersect on or below the $I^R(\bar{Z}^R)$ curve, hence

strictly below the $I^U(\bar{Z}^U)$ curve. So all U -households abstain from pool i . But then our refinement gives $\kappa_i \geq 2/3$, which contradicts $\kappa_i \leq \bar{\kappa}_i < 2/3$.

If $\kappa_i > \bar{\kappa}_i$, then by contributing Q_i to pool i , an R -household can achieve utility higher than $I^R(\bar{Z}^R)$, a contradiction. QED

LEMMA 6. There exists a pool j such that $\lambda_j^U > 0$ and $\lambda_j^R = 0$.

Proof of Lemma 6. This is immediate from Lemma 5 and the fact that $\sum_{j \in J} \lambda_j^U = \sum_{j \in J} \lambda_j^R$. QED

LEMMA 7. (Almost) all U -households achieve the utility I^U .

Proof of Lemma 7. Consider the nonnull set of U -households who act by themselves on pool j (of Lemma 6). Then $\kappa_j = 1/3$. Let X be the consumption of these households. We claim that $X = Z^U$. This follows because I^U (by definition) gives the maximum utility they could achieve on the $1/3$ -price line. So X is on or below I^U . On the other hand, by Lemma 2, X is on or above I^U . It follows that X is on I^U . But, since utilities are strictly concave, $X = Z^U$. Since (almost) all U -households are on the same indifference curve at equilibrium, the lemma follows. QED

Completion of the Proof of Theorem 3.1. Consider j with $Q_j \geq 1$. We shall show that $\kappa_j = 1/3$ and $\lambda_j^R = 0$. By our refinement, $\kappa_j \geq 1/3$. If $\kappa_j > 1/3$, then U -households can achieve utility strictly above I^U (by contributing one unit on pool j), contradicting Lemma 7. This shows that $\kappa_j = 1/3$. From the single crossing property, I^R lies strictly above the $1/3$ -price line. Hence by acting on pool j , reliable households can only obtain utility strictly below their lower bound I^R (see Lemma 4), so $\lambda_j^R = 0$.

Next consider j with $1 > Q_j > Q_{j^*} \equiv \varphi(Z)$. Let \tilde{Z}^U be the intersection of I^U with the Q_j -quantity line, and let $1/3 < \kappa(\tilde{Z}^U) < 2/3$ be the delivery rate corresponding to the dotted line from $(1,0)$ to \tilde{Z}^U . (See Figure V.) We claim that $\kappa_j = \kappa(\tilde{Z}^U)$. If $\kappa_j > \kappa(\tilde{Z}^U)$, then U -households can achieve utility strictly above I^U by contributing Q_j on j , contradicting Lemma 7. If $\kappa_j < \kappa(\tilde{Z}^U)$, then all households obtain at best strictly less than their respective I^U or I^R via pool j , and so (by Lemmas 2 and 4) they will all abstain from acting on j in any perturbation; and they by our refinement $\kappa_j \geq 1 > \kappa(\tilde{Z}^U)$, again a contradiction. This proves that $\kappa_j = \kappa(\tilde{Z}^U)$. Then R -households can at best obtain utility strictly below I^R via pool j , which (by Lemma 4) implies that λ_j^R

= 0. Now if $\lambda_j^U > 0$, we must have $\kappa_j = 1/3 < \kappa(\bar{Z}^U)$, again a contradiction. So $\lambda_j^U = 0$ also.

Now consider j with $Q_j \leq Q_{j^*}$. If $\lambda_j^U > 0$, then

$$\kappa_j = \frac{(1/3)\lambda_j^U + (2/3)\lambda_j^R}{\lambda_j^U + \lambda_j^R} < \frac{2}{3},$$

contradicting Lemma 3. So only R -households can act on these pools. We claim that $\lambda_j^R = 0$ if $j < j^*$. For if $\lambda_j^R > 0$ and $j < j^*$, then $\kappa_j = 2/3$ and the households active on pool j get utility strictly below I^R with the small quantity limit Q_j (recall that they get to I^R on the $2/3$ -price line by means of $Q_{j^*} > Q_j$). This contradicts Lemma 4. Thus, all R -households act on pool j^* alone, and consume $Z^R = Z$ on the I^R indifference curve.

It only remains to verify that the prices κ_j conform to Figures V and VI for all the pools $j < j^*$. This follows from a now-familiar and routine argument. Assume that the refinement is made using a d -external agent, for some fixed $d \geq 1$. Let $\kappa(\bar{Z}^R)$ correspond to the dotted line joining $(1,0)$ to \bar{Z}^R in Figure V (with $\kappa(\bar{Z}^R) < d$). If $\kappa_j < \kappa(\bar{Z}^R)$, all households abstain from j , so refinement gives $\kappa_j \geq d$, a contradiction. If $\kappa_j > \kappa(\bar{Z}^R)$, then R -households can achieve utility above I^R via pool j , again a contradiction. Finally, if the slope of the line from $(1,0)$ to \bar{Z}^R exceeds d , pool j must have $\kappa_j = d$ since all households abstain from j , and the external agent delivers d . QED

XI.B. Proof of Theorem 3.2 (The Case of the Missing Quantity $\varphi(Z) \notin \mathbb{Q}$)

The analysis is very similar to the case $\varphi(Z) \in \mathbb{Q}$, so we shall only give a brief informal summary. First, redefine the region Ω by replacing I^R with the indifference curve through Y (see Figure VIII). This adapts (*) to the missing quantity case. Let us suppose that the grid size is sufficiently small to ensure that the curve $I^R(W)$ is above the line joining $(1,0)$ to W , where $W \equiv$ the intersection of Q_{k+1} -quantity line and I^U (see Figure VIII). Lemmas 1 to 7 hold without change. When the reliable *strictly* prefer Y to W (Case 1, Figure VIII) we get the uniqueness of our separating equilibrium. When they *strictly* prefer W to Y , the splitting equilibrium of Case 2 (Figure IX) is also unique (in delivery rates and the distribution of actions and consumptions). It is only when they are indifferent between W and Y (the "degenerate" scenario) that uniqueness breaks down a *little* bit (but it still holds at many

levels, in particular in terms of the κ_j and the utility levels of equilibrium consumption). For in this scenario, we can take an arbitrary (suitably small) positive measure μ of R -households and mix them with $\tilde{\mu}$ measure of U -households so that $(\mu(2/3) + \tilde{\mu}(1/3))/(\mu + \tilde{\mu}) = \kappa_{k+1}$, where κ_{k+1} , conforms to the slope from $(1,0)$ to W . The μ and $\tilde{\mu}$ populations of R - and U -households join pool $k + 1$ and consume W . The rest of the U - and R -households join pools J and k , and consume Z^U and Y , respectively. Since μ was arbitrary, this gives a continuum of equilibria in terms of consumption. But notice that prices (i.e., the κ_j) as well as utility levels (i.e., I^U and I^R) are still invariant across these equilibria. QED

APPENDIX

A. Proof of Theorem 1

Fix d such that $d \geq \max \{d_s^h : h \in H, s \in S\}$ for all $s \in S$. Let us consider the finite-type continuum (generalized) game Γ_ϵ with a d -external agent who contributes $\epsilon > 0$ on every pool $j \in J$. Then type-symmetric household choices $\varphi \equiv (\varphi^1, \dots, \varphi^H) \in \Sigma \times \dots \times \Sigma$, where $\Sigma = \{\theta \in \mathbb{R}_+^J : \theta_j \leq Q_j\}$, give rise to the delivery rates,

$$K_{sj}^\epsilon(\varphi) = \frac{\epsilon d + \sum_{h \in H} \varphi_j^h d_s^h}{\epsilon + \sum_{h \in H} \varphi_j^h},$$

for $j \in J$ and $s \in S$. For any φ , the feasible set of strategies available to $t \in (h - 1, h]$ is $\Sigma_\epsilon^h(\varphi) \subset \Sigma$, where

$$\Sigma_\epsilon^h(\varphi) = \left\{ \theta \in \Sigma : \chi^h(\theta, K^\epsilon(\varphi)) \equiv e^h - \sum_{j \in J} \theta_j d^j + \sum_{j \in J} \theta_j K_j^s(\varphi) \in \mathbb{R}_+^S \text{ and } \theta_j > 0 \Rightarrow \theta_i = 0 \text{ if } i \neq j \right\}.$$

His best reply in Γ_ϵ is

$$\tilde{\Sigma}_\epsilon^h(\varphi) = \sum_\epsilon^h(\varphi) = \arg \max_{\theta \in \Sigma_\epsilon^h(\varphi)} u^h(\chi^h(\theta, K^\epsilon(\varphi))).$$

By the maximum principle, $\tilde{\Sigma}_\epsilon^h$ is upper semi-continuous and

nonempty valued. Let $\text{Co}(\Sigma_\epsilon^h(\varphi))$ denote the convex hull of $\Sigma_\epsilon^h(\varphi)$. Consider the point-to-set map on $\Sigma \times \cdots \times \Sigma$ given by

$$(\varphi^1, \dots, \varphi^H) \mapsto \text{Co} \left(\sum_\epsilon^1(\varphi) \right) \times \cdots \times \text{Co} \left(\sum_\epsilon^H(\varphi) \right).$$

It can be easily checked that the conditions of Kakutani's theorem are met, so there exists $(\varphi^1(\epsilon), \dots, \varphi^H(\epsilon))$ such that $\varphi^h(\epsilon) \in \text{Co}(\Sigma_\epsilon^h(\varphi))$ for all $h \in \mathcal{H}$. By Caratheodory's theorem, there exist $J + 1$ points $\varphi^{h1}(\epsilon), \dots, \varphi^{h(J+1)}(\epsilon)$ in $\Sigma_\epsilon^h(\varphi)$ and weights $\lambda^{h1}(\epsilon), \dots, \lambda^{h(J+1)}(\epsilon)$ such that $\sum_{j=1}^{J+1} \lambda^{hj}(\epsilon) = 1$ and $\varphi^h(\epsilon) = \sum_{j=1}^{J+1} \lambda^{hj}(\epsilon) \varphi^{hj}(\epsilon)$. Select a subsequence of $\epsilon(n) \rightarrow 0$ such that $\lambda^{hj}(\epsilon(n)) \rightarrow \lambda^{hj}$ and $\varphi^{hj}(\epsilon) \rightarrow \varphi^{hj}$ for all $h \in \mathcal{H}$ and $j = 1, \dots, J + 1$; and (consequently) $K_{sj}^{\epsilon(n)} \rightarrow K_{sj}$ for all $j \in J$ and $s \in S$. Partition $(h - 1, h]$, starting from left to right, into $J + 1$ intervals of length $\lambda^{h1}, \dots, \lambda^{h(J+1)}$ and let the households in the j th-interval choose contributions φ^{hj} and consume $x^{hj} \equiv \chi^h(\varphi^{hj}, K) \in \mathbb{R}_+^S$. It is easy to check that (K, φ, x) is a refined equilibrium with $(K^{\epsilon(n)}, \varphi(\epsilon(n)), x(\epsilon(n)), \epsilon(n))_{n=1}^\infty$ serving as its perturbation sequence, where $\varphi(\epsilon(n))$ and $x(\epsilon(n))$ are defined as follows. Partition each $(h - 1, h]$ from left to right into $J + 1$ intervals of length $\lambda^{h1}(\epsilon(n)), \dots, \lambda^{h(J+1)}(\epsilon(n))$ and for t in the j th interval, put

$$\begin{aligned} \varphi^t(\epsilon(n)) &= \varphi^{hj}(\epsilon(n)) \\ x^t(\epsilon(n)) &= \chi^h(\varphi^{hj}(\epsilon(n)), K^{\epsilon(n)}). \end{aligned}$$

B. The Optimistic External Delivery

Consider our canonical example of insurance, with three states, described in Section V.

Let $E_d = (\kappa, \varphi, x)$ be an equilibrium with external deliveries d , supported by the perturbation $E_d(n) = (\kappa(n), \varphi(n), x(n), \epsilon(n))$. Let $M \equiv \max_h \max_s d_s^h \leq \underline{d} < d$. We claim that there exists $\underline{\kappa}$ such that $E_{\underline{d}} = (\underline{\kappa}, \varphi, x)$ is a refined equilibrium, where $\underline{\kappa}_j = \kappa_j$ for all active $j \in J$, and $\underline{\kappa}_j \leq \kappa_j$ for all $j \in J$. To see why, define

$$\underline{\kappa}_j = \begin{cases} \kappa_j & \text{if } j \text{ is active} \\ \underline{d} & \text{if } j \text{ is inactive and } \kappa_j \geq \underline{d} \\ \kappa_j & \text{if } j \text{ is inactive and } \kappa_j < \underline{d} \end{cases}$$

$$\underline{\kappa}_j(n) = \begin{cases} \kappa_j & \text{if } j \text{ is active} \\ \underline{d} & \text{if } j \text{ is inactive and } \kappa_j \geq \underline{d} \\ \kappa_j(n) & \text{if } j \text{ is inactive and } \kappa_j < \underline{d} \end{cases}.$$

Clearly (κ, φ, x) is an equilibrium, since opportunities are no better for households than at (κ, φ, x) , and the old equilibrium choices are still available.

To construct the perturbation, start with agents acting as in E_d , and make the following changes. Put $1/n$ -external promises delivering \underline{d} (per promise) on each inactive pool with $\kappa_j \geq \underline{d}$. Clearly, the real agents optimize by avoiding pool j , which promises and delivers $\kappa_j(n) = \underline{d} \leq \kappa_j$. For inactive pools j with $\kappa_j < \underline{d} < d$ we know that there is a positive measure $\lambda_j^U(n) + \lambda_j^R(n) > 0$ of real households contributing to pool j in the perturbation $E_d(n)$ (otherwise $\kappa_j(n) = d$, so $\kappa_j = \lim \kappa_j(n) = d > \underline{d}$). Note that, for large enough n ,

$$\frac{\lambda_j^U(n)(1/3) + \lambda_j^R(n)(2/3)}{\lambda_j^U(n) + \lambda_j^R(n)} \leq \kappa_j(n) < \underline{d} < d.$$

It follows that we can find $\lambda_j^O(n)$ promises (of the \underline{d} -external agent) such that

$$\frac{\lambda_j^U(n)(1/3) + \lambda_j^R(n)(2/3) + \lambda_j^O(n) \underline{d}}{\lambda_j^U(n) + \lambda_j^R(n) + \lambda_j^O(n)} = \kappa_j(n).$$

Clearly, $\lambda_j^U(n) \rightarrow 0$, and $\lambda_j^R(n) \rightarrow 0$ as $n \rightarrow \infty$, since $\lambda_j^U(n)$ and $\lambda_j^R(n)$ occur in the perturbation $E_d(n)$. Therefore, $\lambda_j^O(n) \rightarrow 0$ as $n \rightarrow \infty$ otherwise, from the above display, $\kappa_j(n) \rightarrow \underline{d}$, contradicting that $\kappa_j(n) \rightarrow \kappa_j < \underline{d}$. So, in the perturbation $E_d(n)$, let $\lambda_j^U(n)$ and $\lambda_j^R(n)$ measures of U and R households contribute on j , and let the external agent contribute $\lambda_j^O(n)$.

In the original perturbation $E_d(n)$, all households are indifferent between the pools on which their class acts. It follows that households of class $h \in \{U, R\}$ are optimizing in our perturbations. QED

STATE UNIVERSITY OF NEW YORK AT STONY BROOK
YALE UNIVERSITY

REFERENCES

- Cho, In-Koo, and David M. Kreps, "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, CII (1987), 179-221.
- Dubey, Pradeep, and John Geanakoplos, "Insurance Contracts Designed by Competitive Pooling," Cowles Foundation Discussion Paper No. 1315, 2001.
- Dubey, Pradeep, John Geanakoplos, and Martin Shubik, "Default and Punishment in General Equilibrium with Incomplete Markets," Cowles Foundation Discussion Paper No. 1304, 2001.
- Rothschild, Michael, and Joseph E. Stiglitz, "Equilibrium in Competitive Insur-

- ance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, C (1976), 629-650.
- Selten, Reinhard, "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, IV (1975), 25-55.
- Spence, Michael, "Job Market Signaling," *Quarterly Journal of Economics*, LXXXVII (1973), 355-374.