On Rothschild-Stiglitz as Competitive Pooling*

Alberto Martin†

February 2006

Abstract

Dubey and Geanakoplos [2002] have developed a theory of competitive pooling, which incorporates adverse selection and signaling into general equilibrium. By recasting the Rothschild-Stiglitz model of insurance in this framework, they find that a separating equilibrium always exists and is unique.

We prove that their uniqueness result is not a consequence of the framework, but rather of their definition of refined equilibria. When other types of perturbations are used, the model allows for many pooling allocations to be supported as such: in particular, this is the case for pooling allocations that Pareto dominate the separating equilibrium.

Keywords: competitive pooling, insurance, adverse selection, signalling, refined equilibrium, separating equilibrium

JEL Classification: D4, D5, D41, D52, D81, D82

*I am deeply indebted to Paolo Siconolfi for numerous discussions and encouragement. I a particularly grateful to Max Amarante and John Geanakoplos for discussions and important insights. I would also like to thank, without implicating, Filippo Taddei, Guido Sandleris, Wouter Vergote and two anonymous referees. I acknowledge financial support from the Spanish Ministry of Education and Science (grant SEJ2005-01126), from the Generalitat de Catalunya, DURSI and from CREA-Barcelona economics.

†CREI, Universitat Pompeu Fabra, 25-27 Ramon Trias Fargas, Barcelona, Spain (08005) Email: alberto.martin@upf.edu

1 Introduction

The goal of this paper is to show that, when the Rothschild and Stiglitz [7] model of insurance is recast as a model of competitive pooling, it admits many equilibria. This is contrary to the recent result by Dubey and Geanakoplos [1], by which the model has a unique equilibrium that coincides with the separating equilibrium of Rothschild and Stiglitz. It is shown that this uniqueness result is not a consequence of the competitive pooling framework per se, but stems rather from the particular tremble that is used in defining refined equilibria.

In their paper, Dubey and Geanakoplos analyze the phenomenon of pooling in a general equilibrium framework which incorporates adverse selection and signaling. Their setting is characterized by households that differ in their reliability and trade by making contributions to a collection of pools: these, in turn, differ through their quantity limits on contributions, so that the maximum amount of "shares" that any given household can buy differs across pools. In equilibrium, pools with lower quantity limits sell for a higher price, even though each household's deliveries are the same at all pools.

One of the main features of their construction is that they use a tremble "on the market" in order to specify equilibria: since pools that are inactive at equilibrium receive no contributions and make no payments, the anticipated returns of these pools cannot be pinned down by rational expectations. To overcome this problem, they perturb the equilibrium by introducing an external agent that makes contributions to all pools along the perturbation, constructing in this way a refined equilibrium.

After developing the framework, Dubey and Geanakoplos reinterpret Rothschild and Stiglitz's model of insurance (henceforth, [RS]) as a particular case of competitive pooling. They show that, by recasting the model in this perfectly competitive fashion, an equilibrium always exists under the assumptions of exclusivity and no cross-subsidization among the pools. Furthermore, the equilibrium is proved to be unique and to correspond to the [RS] separating equilibrium.

The result of Dubey and Geanakoplos regarding existence, although in stark contrast to the features of the strategic model, is based on fixed point arguments that are common in the literature. Their uniqueness result, however, is more surprising at first glance. The reason is that, typically, models with incomplete information have many different equilibrium allocations supported by different beliefs about inactive markets.

The object of this paper is precisely to show that the uniqueness result of Dubey and Geanakoplos is not a consequence of the competitive pooling framework per se, but stems rather from the particular tremble they use in defining refined equilibria. In this regard, they use what they define as an optimistic tremble, by which the external agent is assumed to deliver more than any of the households in the economy. We prove that under different trembles, all incentive compatible and feasible pooling allocations that Pareto dominate the [RS] separating equilibrium can be supported as refined equilibria of the model. On the other hand, the model has a unique separating equilibrium regardless of the tremble that is assumed.

In games with asymmetric information, it is well known that the set of equilibria generally depends on the refinement of Nash equilibrium that is used. In the particular case of adverse selection, Hellwig [5] has also pointed out that the set of equilibria depends on the exact manner in which competition is modeled, or - more generally - on the game form that specifies the interaction between informed and uninformed agents in the model. It can thus be thought that the characteristics of the tremble invoked represent to the model of competitive pooling what the details of game theoretic interaction and the use of equilibrium refinements are to strategic models of adverse selection: one must fully specify them in order to discuss the properties of equilibria.

A similar point has been previously made by Gale ([2],[3],[4]) in a series of papers in which he analyzes markets with adverse selection from a general equilibrium perspective. In Gale [2], he discusses the use of different refinements in order to specify equilibria of the competitive model with adverse selection: in particular, he shows that separating equilibria are the most robust, in the sense that they are the only ones to satisfy the Kohlberg-Mertens [6] stability criterion. Although we do not explore this issue for lack of space, the working version of the paper argues that this is precisely the underlying feature that drives the uniqueness result of Dubey and Geanakoplos: by assuming a very high delivery on behalf of the external agent, they are implicitly using the Kohlberg-Mertens stability criterion to define refined equilibria. By reducing the contribution of the external agent, we are in fact weakening the refinement concept to resemble the original notion of "trembling-hand".

The structure of the paper is as follows: Section 2 describes the framework used by Dubey and Geanakoplos. Section 3 identifies the set of feasible, incentive compatible allocations of the model and proves that the latter has many pooling equilibria and a unique separating equilibrium.

2 Rothschild-Stiglitz as competitive pooling

2.1 Households

Consider a one-good economy with a continuum of households denoted by a superscript $t \in (0, 1]$. Households are divided into two types, "reliable" and "unreliable", with population measures of λ_R and λ_U, respectively. All uncertainty in this economy is idiosyncratic, and each of the $S = n$ states of nature can be of two kinds for any one household: a "good" state, in which endowment is positive (and, to simplify, equal to one unit), or a "bad" state, in which endowment is zero. A reliable household is defined as having zero endowment in r out of the total n states, whereas an unreliable household has a bad state in u of them, where $u > r$. If consumption of household t in state s is denoted by x_s^t, the utility function common to all households is given by,

$$U(x_1^t, x_2^t, ..., x_n^t) = \frac{1}{n} \sum_{s=1}^{n} u(x_s^t), \text{ for all } t \in (0, H] \tag{1}$$

where it is assumed that $u' > 0$, $u'' < 0$ and $\lim_{x \to 0} u'(x) = \infty$.

This setup looks very similar to that of the [RS] model of insurance. There is, however, an important difference that arises from explicitly considering the possible states of the world. In this sense, although all reliable households have zero endowment in r out of the n states, the *actual* states which are bad differ across reliable households. This characteristic of the model requires some care in the way in which households and types are defined: however, we avoid dealing with it directly by making the following assumption

(A1) Given any non-null set of households of type R (U), a proportion $\frac{r}{n}$ ($\frac{u}{n}$) of them will have a bad state in any given state.

(A1) will be assumed to hold not just in equilibrium but also in deviation from equilibrium. Although this assumption might seem to raise the usual difficulties, it is of widespread use in the literature and one of the original assumptions of [RS]. In any case, it is possible to show that (A1) can be derived from a fairly standard general equilibrium setting with contingent endowment.

2.2 Trade

It is assumed that households can only trade through contributions to a collection of pools indexed by $j \in J = \{1, ..., J\}$: by contributing to pools, households acquire the right to a share of the latter's resources but also the obligation to make certain deliveries. Pools differ in the maximum amount that can be contributed to them, which we denote by the vector $Q \in R^J_+$: as a convention, we assume $Q_1 < Q_2 < ...1 < .. < Q_J$.

Consequently, each household is entitled to contribute $0 \leq \varphi^t_j \leq Q_j$ promises to pool j, which oblige him to make state-contingent deliveries $d^t \in R^s_+$ per unit of promise: without loss of generality, we assume $d^t = e^t$ for all t. Additionally, the contribution entitles the household to a share of the pool's total resources. We maintain the following assumption throughout the paper,

(A2) No cross-subsidization: the payments made by each pool must necessarily come from their own revenues.

Therefore, the economy is defined by $((u^t, e^t, d^t)_{t\in(0,1]}, (Q_j)_{j\in J})$. When a household contributes φ^t_j to pool j, he agrees to delivering φ^t_j in his good states and receiving $\varphi^t_j K_{js}$ from the pool in each state $s \in S$, where by (A2) $K_{js}(\varphi) = \frac{1}{\bar{\varphi}_j}\int_0^1 \varphi^t_j d^t_s dt$ denotes the return of pool j in state s, $\bar{\varphi}_j = \int_0^1 \varphi^t_j dt$ denotes the pool's aggregate holding of promises and $\varphi \in R^{J\times(0,1]}_+$ denotes the contributions of all households across pools. Note that no single household can affect the return of any pool j, which is denoted by the vector $K_j = (K_{j1}, ..., K_{jS}) \in R^S_+$. From their perspective then, the feasible trades are given by the fixed matrix $K = K(\varphi) \in R^{J\times S}_+$, in which the (j, s)th element is given by $K_{js}(\varphi)$ as previously defined. The consumption of household $t \in (0, 1]$ in each state s, when his contributions to pools are represented by the vector φ^t is given by,

$$x^t_s = \chi^t_s(\varphi^t, K) = e^t_s + (K_s - d^t_s)\varphi^t \tag{2}$$

where $K_s = (K_{1s}, ..., K_{Js}) \in R^J_+$ and $\varphi^t \leq Q \in R^J_+$. His budget set can then be expressed as,

$$\beta^t(K) = \{(\theta, y) \in R^J_+ \times R^S_+ : y = \chi^t(\theta, K)\} \tag{3}$$

Given this setup, we are ready to define an equilibrium of the economy. The only thing that needs to be addressed before doing so refers to the specification of the return vector K_j for pools that are inactive at equilibrium. Since these pools receive no contributions and make no payments, their returns are not specified and cannot be anticipated by rational households. However, a complete specification of the latter's optimization problem must include some beliefs about the return of inactive pools, which must therefore be specified in order to obtain a characterization of equilibria.

2.3 The equilibrium refinement

Dubey and Geanakoplos deal with the anticipated deliveries of inactive pools by establishing an analogy between them and game theoretic beliefs "off the equilibrium path". Building on this analogy, and recalling Selten's use of trembles to deal with the latter problem, they too invoke a tremble but "on the market" instead of on households.

In this regard, they consider an external ε-agent who - along a sequence - contributes nonnegative amounts $\varepsilon(n) = (\varepsilon_j(n))_{j\in J} \in R_+^J$ to every pool and delivers a fixed vector $d = (d, ..., d) \in R_+^S$ per unit contributed to any particular pool. They then require that $d \geq \max_{h=R,U} d_s^h$ for all $s \in S$, and call any d satisfying this requirement *optimistic*. Finally, they assume that $\varepsilon(n) \to 0$ as $n \to \infty$, suggesting that the agent might be interpreted as a government that guarantees delivery on the first infinitesimal promises.[1] In formal terms, they define a refined equilibrium as follows:

Definition 1 *An equilibrium* $E = (K, \varphi, x) \in R_+^{J\times S} \times R_+^{J\times(0,1]} \times R_+^{S\times(0,1]}$ *of the economy* $((u^t, e^t, d^t)_{t\in(0,1]}, (Q_j)_{j\in J})$ *is said to be a refined equilibrium if there is a sequence* $E_d = (K(n), \varphi(n), x(n), \varepsilon(n)) \in R_+^{J\times S} \times R_+^{J\times(0,1]} \times R_+^{S\times(0,1]} \times R_+^J$ *such that* d *is optimistic,* $\varphi(n)$ *and* $x(n)$ *are measurable for all* $n = 1, 2, 3, ...$ *and*

(1) $\varepsilon(n) \to 0,\ K(n) \to K$ *and* $\varphi^t(n) \to \varphi^t$*,* $x^t(n) \to x^t$ *for almost all* t

(2) $(\varphi^t(n), x^t(n)) \in \arg\max_{(\theta,y)\in\beta^t(K(n))} u^t(x)$ *for almost all* t *and all* n

(3) $\varepsilon_j(n) > 0$ *if* $\bar{\varphi}_j(n) = 0$*, for all* $j \in J$ *and all* n*,*

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

$K_{js}(n) = \frac{1}{\varepsilon_j(n)+\bar{\varphi}_j(n)}[\varepsilon_j(n)d + \int_0^H \varphi_j^t(n) d_s^t dt]$

A refined equilibrium is then defined as the limit of a sequence of equilibria of the perturbed economy, along which the external agent contributes to inactive pools and hence allows for the household optimization problem to be well defined. Under our assumptions, it can be proved that a refined equilibrium always exists in the finite type continuum model with different pools.

Before proceeding, it is worthwhile to highlight some aspects of the framework. Note that, in principle, a pool's return could vary across states: however, due to (A1) and to the fact that the external agent delivers the same across all states, $K_{sj}(n)$ will be fully determined and state-invariant for a given composition of households that contribute to pool j. Hence, the optimization

[1] To simplify computations, they actually drop the contributions of the external agent in those pools that are active in equilibrium.

problem will be identical for all households of the same type, and we can therefore rewrite the model in these terms.

We also follow Dubey and Geanakoplos by maintaining the following assumption in what remains of the paper:

(A3) Exclusivity constraint: each household can contribute to at most one pool[2]

Hence, we can define $1 - p_R = \frac{r}{n}$ and $1 - p_U = \frac{u}{n}$ and rewrite (1) as

$$u^h = p_h u(x_G) + (1 - p_h)u(x_B) \tag{4}$$

where, $h = R, U$ and

$$\begin{aligned} x_G &= (1 - \varphi_j^h) + \varphi_j^h K_j \\ x_B &= \varphi_j^h K_j \end{aligned}$$

where, because of (A3), φ_j^h represents the only nonzero element of the vector of contributions φ^h. Additionally, and with some abuse of notation, we henceforth let $K_j \in R_+$ represent the return of pool j, which for the reasons outlined above is invariant across states. Finally, we use $\bar{p} = \lambda_R p_R + \lambda_U p_U$ to denote the population's average probability of a good state. This representation of the model, in terms of good and bad states, is has the advantage of allowing the use of the traditional graphical representation of [RS].

Note that if the return of pool j across states is constant and given by K_j, a household that contributes φ_j to it gives up φ_j consumption in the good states and receives $K_j \varphi_j$ in every state. On the net, then, such a household foregoes $(1 - K_j)\varphi_j$ units of consumption in his good states in exchange for $K_j \varphi_j$ units in the bad states, thereby shifting consumption from the former to the latter at a rate of $\frac{K_j}{1-K_j}$.

3 Equilibria of the Model

The present section analyzes how the set of equilibria of the model depends on the external agent's delivery rate. We prove that, besides the [RS] separating equilibrium, the model also displays

[2] This assumption, which delivers a non-convex budget set, does not affect the existence of an equilibrium, as Dubey and Geanakoplos show.

many pooling equilibria: in particular, all pooling allocations that Pareto dominate the former can be supported as such. It must be remembered that (A1), (A2) and (A3) are assumed to hold throughout the analysis. We proceed by identifying the set of feasible, incentive compatible allocations of the model and then analyze which one of them can be supported as pooling or separating refined equilibria.

3.1 Feasible, incentive compatible allocations

The following lemma identifies the set of feasible, incentive compatible separating and pooling allocations of the model. An allocation is said to be feasible if, in all open pools, the payments distributed among contributing households equal the revenues received from household deliveries. Regarding incentive compatibility, an allocation is said to be incentive compatible if households do not gain by deviating to other open pools or by changing the level of their contribution to any open pool.

It is obviously the case that in separating allocations, all but two pools will be closed, while all but one will be closed in pooling allocations. Up to now, it has been assumed that there is a finite number of pools. For the remaining of the paper, however, we will consider the competitive insurance model with a continuum of pools $j \in [0, J]$, where $Q_J > 1$ and $j < j^* \iff Q_j < Q_{j^*}$.[3]

Lemma 1 *There exist a continuum of feasible, incentive compatible separating and pooling allocations.*

Proof. See Appendix. ■

The intuition behind the previous lemma can be explained as follows. In the case of separating allocations, whenever unreliable types contribute to a pool $j_U \leq 1$, there is a unique pool $j_R < j_U$ such that reliable (unreliable) types contributing to the former (latter) is both feasible and incentive compatible. In the case of pooling allocations, any allocation in which both types contribute to the same pool j an amount equal to Q_j is also feasible and incentive compatible.

This identifies the set of feasible, incentive compatible allocations of the competitive insurance model. We will now identify the subset of such allocations that can be supported as refined equilibria, and assess how this subset changes depending on the delivery rate of the external agent. Before proceeding, however, let us consider the case of $d = 0$.

[3] The use of a continuum of pools poses some potential problems for existence in the model: however, it is done here solely to avoid some uninteresting complications and in order to simplify the exposition. Qualitatively, the results would be the same if there were finitely many pools.

Remark 2 *All of the allocations mentioned in Lemma 1 can be supported as refined equilibria if the external agent delivery rate is set at $d = 0$.*

The previous remark merits no further comments, since it's clear that agents will choose only among the pools that are active in equilibrium if they expect all inactive ones to deliver nothing. However, since all the allocations considered in the lemma are incentive compatible and feasible, this is equivalent to saying that they can be supported as refined equilibria. Thus, the model displays many refined equilibria when $d = 0$.

3.2 Pooling Equilibria

The present subsection shows that, given an external agent delivery rate $d \in [p_U, \bar{p}]$, any feasible and incentive compatible pooling allocation of the model can be supported as a refined equilibrium *provided that the unreliable households prefer it to the [RS] separating allocation.* In particular, this means that when such a pooling allocation Pareto dominates the [RS] separating one, it can be supported as a refined equilibrium if the delivery rate of the external agent lies in the interval specified above. The statement of the proposition, which is proved in the Appendix, is as follows:

Proposition 1 *If the external agent's delivery rate is set at $d \in [p_U, \bar{p}]$, all of the feasible, incentive compatible pooling allocations of the competitive insurance model can be supported as refined equilibria as long as the unreliable households prefer it to the Rothschild-Stiglitz separating allocation.*

Since the proof of the proposition is constructive and thus somewhat cumbersome, it is worthwhile to portray the intuition behind it. The challenge is to construct a perturbation by switching disjoint sets of households into pools that are inactive at equilibrium and letting the external agent contribute to them as well: additionally, this must be done in such a way as to avoid deviations along the perturbation.

We achieve this by moving unreliable households to pools with higher quantity constraints than the pooling allocation (i.e., pools "to the left" of the pooling allocation in Figure 3) and reliable households to those with lower constraints (pools "to the right"). The measures of reliable and unreliable households that are moved are chosen so as to keep the return of the pooling allocation constant. Once households have been moved in this fashion, the external agent is introduced in order to make unreliable households indifferent among all the pools to which their type contributes, while the same must be true for reliable households: note that the latter is possible only because the external agent's delivery is weakly below the average of the economy. Additionally, it must be considered that this construction does not rule out the possibility of having some pools where the external agent is the sole contributor: it must be the case, however, that these pools are not preferred by any household to the pooling allocation. Any perturbation constructed along these lines will clearly satisfy the definition of refined equilibria, and will therefore support the pooling allocation as such.

The result by which pooling allocations may be supported as equilibria for some delivery rates is not surprising, since the main problem for doing so arises from their fragility to deviations by reliable households: it is enough for a pool immediately to the "right" of the equilibrium to exhibit a return higher than $\bar{p}$ in order for it to attract reliable households, eventually destroying the pooling allocation. Thus, such an equilibrium can only be supported by perturbations in which $d \leq \bar{p}$. This is the reason why Dubey and Geanakoplos, by assuming $d = 1$, conclude that the model has no pooling equilibria.

3.3 Separating Equilibria

In the case of separating equilibria, we prove that the model has only one for all delivery rates $d \geq p_U$: the [RS] separating allocation. Considered jointly, the previous and the following propositions show that for delivery rates $d \in [p_U, \bar{p}]$ there are many pooling allocations that can be supported as refined equilibria of the model, whereas only a unique separating allocation - the [RS] separating allocation - can be supported as such for all delivery rates $d \geq p_U$.

Proposition 2 *For all delivery rates $d \geq p_U$, the only separating allocations that can be supported as a refined equilibrium of the competitive insurance model is the Rothschild-Stiglitz separating allocation.*

Proof. We first argue that for any delivery rate $d \geq p_U$, no feasible, incentive compatible separating allocation different from the [RS] one can be supported as a refined equilibrium. To address this, consider first the set of such allocations in which unreliable households are contributing less than they would like to at a return of p_U. For any $d \geq p_U$, however, p_U is the lower bound for anticipated returns of inactive pools, so unreliable households will immediately move their contributions to pool j for which $Q_j = 1$. Alternatively, consider a feasible, incentive compatible separating allocation in which unreliable households contribute to a pool Q_j for $j > 1$. In such a scenario, their contribution to pool j will be equal to one, and so such an allocation will be *de* *facto* equivalent to the [RS] separating allocation. Hence, no other separating allocation can be supported as a refined equilibrium.

It remains then to be shown that the [RS] allocation itself can be supported as a refined equilibrium for any $d \geq p_U$. This must in fact be the case since - as Dubey and Geanakoplos show - it can be supported as such for $d = 1$. Since the effect of lowering d below one is essentially to decrease the return on inactive pools, the [RS] separating allocation must still be a refined equilibrium. If this were not the case, it means that it is possible to make households indifferent between the equilibrium allocation and any deviation to inactive pools when $d = 1$ but not when $d < 1$, a contradiction since inactive pools are less attractive under the latter perturbation. ■

4 Concluding Remarks

The object of this paper has been to show that the set of equilibria in the competitive version of the Rothschild-Stiglitz model of insurance depends on the perturbation that is used. For external agent delivery rates $d \in [p_U, \bar{p}]$, the model allows for many pooling equilibria besides the Rothschild-Stiglitz separating equilibrium: for higher delivery rates, only the latter survives.

The working version of the paper also shows that different delivery rates of the external agent amount to invoking different concepts of equilibrium refinement. Thus, our competitive analysis has a strong analogy in the case of games with asymmetric information, where the set of equilibria depends on the competitive refinement of Nash equilibrium or on the particular game form that is used. Regarding the latter, Hellwig [5] has stressed that the conclusions of strategic models of adverse selection are very sensitive to the details of the game form that is assumed: in the particular case of the Rothschild-Stiglitz model, for example, the existence of equilibrium and its characteristics are deeply affected by whether the game is modeled as having two or three stages. In much the same way, the perturbation used in the competitive version of the model has a substantial effect on the conclusions reached: it alone will determine whether the equilibrium may be pooling or separating and, consequently, whether it will be Pareto optimal or not.

Finally, what type of refinement or external agent delivery should then be used? Although an in-depth discussion of the issue exceeds the scope of this paper, it can nonetheless be said that the right approach should depend on the interpretation that is given to the perturbation. If the perturbation is to be interpreted as households' irrationality or as mistakes which grow vanishingly small at the equilibrium, then it seems reasonable to prioritize robustness in the solution concept:

any other choice would imply an assessment regarding the greater likelihood of some mistakes with respect to others, which seems rather arbitrary at best. On the other hand, if the perturbation is to be interpreted as some type of government signal or action, as suggested by Dubey and Geanakoplos, then there is no *a priori* reason to justify any one delivery rate over another and - even though it should still be taken into consideration - robustness would not seem to be an issue of great concern.[4] In fact, the findings of this paper are very intuitive when applied to the latter setting: if a government wishes to support a pooling equilibrium because of its greater efficiency, then it should "encourage" reliable households to stay there by inducing relatively low anticipated returns in other markets. On the other hand, it is obvious that if the government somehow increases anticipated returns in inactive markets (as the external agent with high delivery rates does) then a pooling allocation cannot survive as an equilibrium.

[4]Note that robustness is still a source of concern because we have shown only that it is possible for the optimal pooling allocation to be a refined equilibrium under certain government delivery rates, not that it will necessarily arise as the equilibrium. This is due to the fact that we have restricted the external agent´s delivery rate to be constant across pools: if this restriction were relaxed, so that the external agent could vary its delivery rate across pools, then it would be possible to construct perturbations where the optimal pooling allocation is the only refined equilibrium of the economy.

References

[1] Dubey, P. and Geanakoplos J., Competitive Pooling: Rothschild Stiglitz Reconsidered, The Quarterly Journal of Economics **117**, 1529-1570 (2002).

[2] Gale, D., A Walrasian Theory of Markets with Adverse Selection, The Review of Economic Studies **59**, 229-255 (1992).

[3] Gale, D., Equilibria and Pareto Optima of Markets with Adverse Selection, Economic Theory **7**, 207-235 (1996).

[4] Gale, D., Signaling in Markets with Two-Sided Adverse Selection, Economic Theory **18**, 391-414 (2001).

[5] Hellwig, M., Some Recent Developments in the Theory of Competition in Markets with Adverse Selection, European Economic Review **31**, 319-325 (1987).

[6] Kohlberg, E. and Mertens, J.F., On the Strategic Stability of Equilibria, Econometrica **54**, 1003-1038 (1986).

[7] Rothschild M., Stiglitz J., Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information, The Quarterly Journal of Economics **90**, 629-649 (1976).

[8] Selten, R., Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games, International Journal of Game Theory **4**, 25-55 (1975).

5 Appendix

Lemma 1. Existence of a continuum of feasible, incentive compatible separating allocations: Consider the Rothschild-Stiglitz separating allocation, where unreliable households contribute $\varphi(x^U) = 1$ to pool $j \in J$ for which $Q_j = 1$. Now consider that all the unreliable households contribute to pool $\bar{j} < j$ instead, so that they choose a bundle x^U along the p_U-price line for which $\varphi(x^U) < 1$ (see Figure 1 below). Define x^R to be the intersection of the $I^U(x^U)$ indifference curve with the p_R-price line, and define $j^* \in J$ so that $Q_{j^*} = \varphi(x^R)$. Clearly, any allocation (x^U, x^R) constructed in this way, in which unreliable households contribute to pool $\bar{j}$ and reliable households to pool j^*, is incentive compatible. Feasibility is obviously guaranteed by the nature of the pools considered, which deliver no more than what they receive. There exist a continuum of allocations that can be constructed in this way.

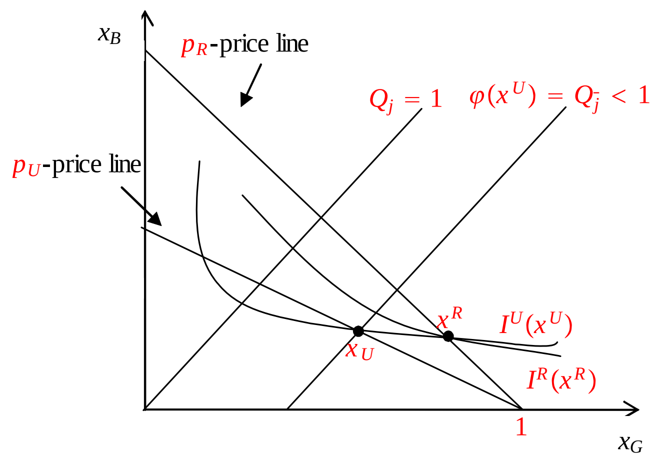

Figure 1: A feasible, IC, separating allocation

Existence of a continuum of feasible, incentive compatible pooling allocations: Define $\bar{x}$ as the point on the $\bar{p}$-price line that satisfies,

$$\varphi(\bar{x}) = \arg\max_{\varphi}[p_R u(1 - \varphi(\bar{x}) \cdot (1 - \bar{p})) + (1 - p_R)u(\varphi(\bar{x}) \cdot \bar{p})]$$

i.e., the tangency point for the reliable households on the $\bar{p}$-price line (see Figure 2 below). Clearly, if j^* is defined so that $Q_{j^*} = \varphi(\bar{x})$, the allocation where both types contribute $\varphi(\bar{x})$ to pool j^* is feasible and incentive compatible. The incentive compatibility of this allocation comes from the fact that reliable households are contributing exactly what they want to the pool: unreliable households, on the other hand, would like to contribute more but are constrained by the quantity limit. Therefore, all households end up contributing the same amount. It follows immediately that a continuum of such allocations can be constructed: in fact, we may do so for all pools $j < j^*$, to which all households will be contributing Q_j for a return of $\bar{p}$.

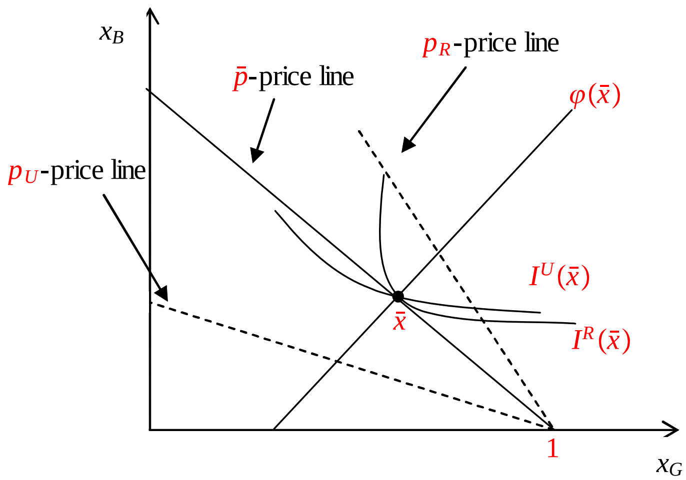

Figure 2: Pooling allocation preferred by reliable types

■

Proposition 1. Since the proof uses graphs to simplify the exposition, let us comment on the notation used. We do so through Figure 3, which depicts the pooling allocation mostly preferred by reliable types. Given the bundle $\bar{x} = (\bar{x}_G, \bar{x}_B)$, we will define $\varphi(\bar{x})$ to be the quantity line passing through $\bar{x}$. In the same manner, we will define $K(\bar{x})$ to be the price line passing through $\bar{x}$ (which, in the case of a pooling allocation, coincides with the $\bar{p}$-price line) and $I^R(\bar{x})$ and $I^U(\bar{x})$ to be, respectively, the reliable and unreliable households' indifference curves passing through $\bar{x}$.

Let all households contribute $\varphi(\bar{x})$ to pool $j^* < J$, where $\varphi(\bar{x}) = Q_{j^*}$ and $\bar{x}$ is the point on the $\bar{p}$-price line for which $\frac{p_R}{(1-p_R)} \cdot \frac{u'(\bar{x}_G)}{u'(\bar{x}_B)} = \frac{\bar{p}}{1-\bar{p}}$. This situation is depicted in Figure 3 above. The choice of Q_{j^*} is motivated to simplify the exposition, but the proof goes through for any other feasible, incentive compatible pooling allocation (i.e., in which only one pool j is active, for some $j < j^*$).It is assumed that the utility of the unreliable households in this pooling allocation is higher than what they would obtain in the [RS] separating allocation. Once it has been proved that the pooling allocation can be supported as a refined equilibrium, it will become clear why this condition is necessary when a reasonably pessimistic tremble is used.

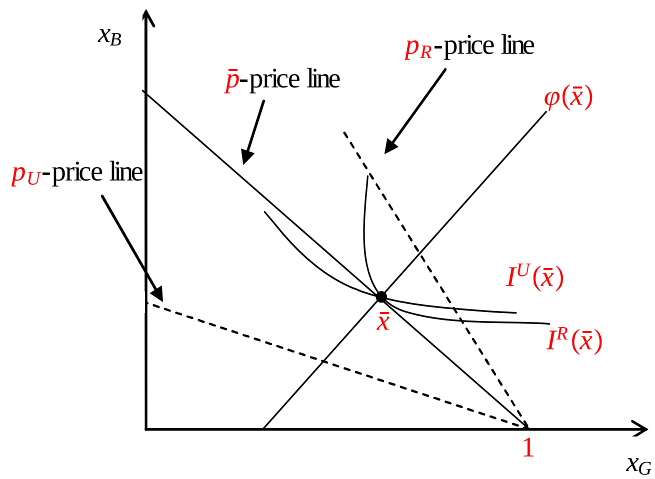

Figure 3: Pooling allocation preferred by reliable types

In order to prove that such a pooling allocation can be supported as a refined equilibrium, all inactive pools must be priced in a way that satisfies the equilibrium refinement. In a first step, we will determine the required prices of inactive pools. Once this has been done, we will construct a sequence that supports the allocation as a refined equilibrium.

Step 1: For pools j with $j > j^*$, the Q_j-quantity line intersects $I^U(\bar{x})$ before it intersects $I^R(\bar{x})$. Let $\hat{Q}^U$ be the quantity whose line intersects $I^U(\bar{x})$ at $\hat{x}^U$, defined in the following way,

$$\hat{x}^U = \arg\min k(x) \tag{5}$$

$$\text{s.t. } I^U(x) = I^U(\bar{x})$$

Then, there are two possibilities:

a) $Q_j > \hat{Q}^U > Q_{j^*}$. In this case, let $K_j \leq K(\hat{x}^U)$ (Region A in Figure 4).

b) $\hat{Q}^U \geq Q_j > Q_{j^*}$, so that the Q_j-quantity line intersects $I^U(\bar{x})$ at $\tilde{x}^U$. In this case, set $K_j \leq K(\tilde{x}^U)$, where the equality is in accordance with the dotted line in Figure 4 that connects (1,0) with $\tilde{x}^U$ (Region B in Figure 4).

For pools j with $j < j^*$, the Q_j-quantity line intersects $I^R(\bar{x})$ before it intersects $I^U(\bar{x})$. Let $\hat{x}^R$ be the point where $I^U(\bar{x})$ intersects the p_R-price line, and let $\hat{Q}^R$ be the quantity whose line passes through it. Once again, there are two possibilities:

a) $Q_{j^*} \geq Q_j > \hat{Q}^R$, so that the Q-quantity line intersects $I^R(\bar{x})$ at $\tilde{x}^R$. In this case, set $K_j \leq K(\tilde{x}^R)$, where the equality is in accordance with the dotted line in Figure 4 that connects (1,0) with $\tilde{x}^R$ (Region C in Figure 4).

b) $Q_{j^*} > \hat{Q}^R \geq Q_j$, in which case $K_j \leq p_R$ (Region D in Figure 4).

Given these returns for all $j \neq j^*$, it can be easily verified that no household can improve his utility by using pool j. Thus, such returns could define an equilibrium: all that needs to be done in order to conclude the proof is to construct a perturbation $E = (K(n), \varphi(n), x(n), \varepsilon(n))$ to show that this equilibrium is refined.

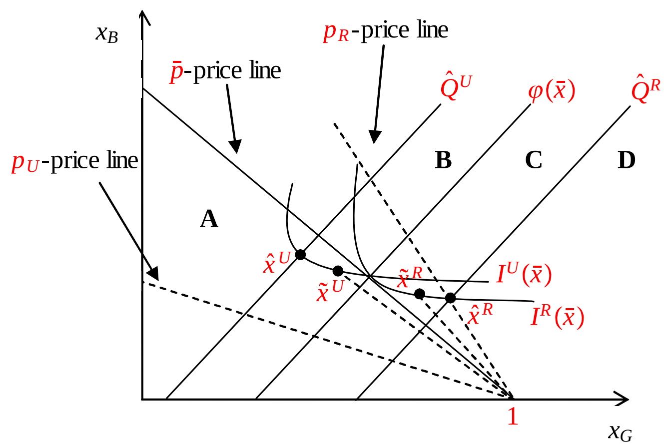

Figure 4: Pools in Regions A, B, C and D

Step 2: The perturbation is constructed by switching disjoint sets of households from their equilibrium actions onto inactive pools, thus defining $\varphi_j^h(n)$ for $j \neq j^*$ and $h = R, U$. In order to calculate the returns of these pools along the ε-sequence, the contributions of the external agent must also be taken into account. Thus, the return of a pool j is, for all $j \neq j^*$and $n = 1, 2, 3...$

$$K_j(n) = \frac{m_j^h(n)\varphi_j(n)p_h + \varepsilon_j(n)d}{m_j^h(n)\varphi_j(n) + \varepsilon_j(n)}$$

where $m_j^h(n)$ is the measure of households of type h moved to pool j along the perturbation. As long as it is possible to construct returns $K_j(n)$ that are consistent with the equilibrium returns considered in Step 1, the pooling allocation can be supported as a refined equilibrium.

It must be noted that there are two possibilities when $d \in [p_U, \bar{p}]$. In the first case, suppose that $d \leq K(\hat{x}^U)$, where $K(\hat{x}^U)$ is defined as in (5) above. With such a low contribution by the external agent, any perturbation which involves not moving any households away from the pooling allocation will suffice: the return of all inactive pools, being solely determined by the external agent, will not be sufficiently high as to induce deviations by any of the households.

Consider instead that the contribution of the external agent d is contained in the interval $(K(\hat{x}^U), \bar{p}]$. In this case, any perturbation will have to imply the movement of disjoint set of households from the pooling allocation to inactive pools. This immediately poses the following questions:

a) What measures of reliable and unreliable households should be moved away from the pooling allocation?

b) How should these measures be distributed among inactive pools?

We answer the first question in such a way as to maintain the return of the pooling allocation unaltered, i.e., we move measures R_n of reliable households and U_n of unreliable ones along the perturbation, where,

$$R_n = U_n \frac{(\bar{p} - p_U)}{(p_R - \bar{p})} \forall n \tag{6}$$

$$R_n \to 0, U_n \to 0 \text{ as } n \to \infty \tag{7}$$

Regarding the distribution of households among inactive pools, it will be done in the following manner. The unreliable households will be uniformly distributed between pools J and j_d, where the latter is the pool whose quantity line passes through x_d^U, the point where $I^U(\bar{x})$ intersects the d-price line. Thus, $\varphi(x_d^U) = Q_{j_d}$ and $K(x_d^U) = d$ (see Figure 5).[5] Thus, each of these pools will receive a measure u_{jn} of such households, where $\int_{j_d}^{J} u_{jn} dj = U_n$ for all n. As for the remaining inactive pools, they won't receive any unreliable households.

With respect to reliable households, we will distribute measures R_n of them - satisfying condition (6) - uniformly among pools between Q_{j^*} and $\hat{Q}^R$.[6] No households will thus be switched to pools j for $Q_j < \hat{Q}^R$.

We shall now prove that a perturbation constructed along the previous lines can support the pooling allocation as a refined equilibrium.

[5] In terms of Figure 5, this means that unreliable households will be distributed among all pools in region A and part of the ones in region B.

[6] In terms of Figure 5, this means that reliable households will be uniformly distributed among all pools in region C.

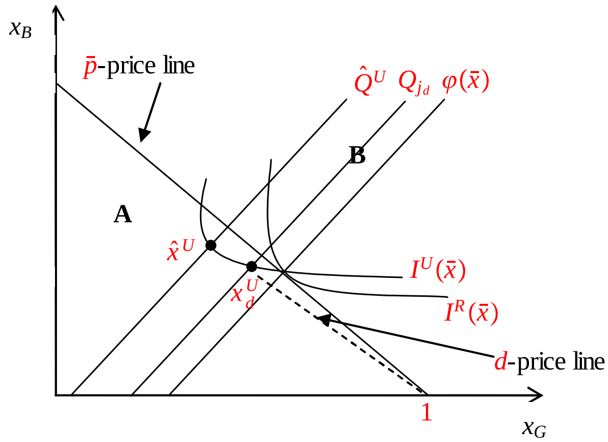

Figure 5: d-price line and corresponding x_d^U

Consider the pools j for which $Q_j > \hat{Q}^U$ (Region A in Figure 4). Let a set of unreliable households of measure u_{jn} move out of pool j^* and contribute to pool j. If $K_j = K(\hat{x}^U)$ for all these pools, unreliable households will always optimize by contributing $\hat{Q}^U$ to any of them. Because $K(\hat{x}^U) < \bar{p}$ and $d \in (K(\hat{x}^U), \bar{p}]$, it is always possible to choose $\varepsilon_j(n) > 0$ to satisfy

$$\frac{u_{jn}\hat{Q}^U p_U + \varepsilon_j(n)d}{u_{jn}\hat{Q}^U + \varepsilon_j(n)} = K(\hat{x}^U)$$

Consider now the pools j for which $\hat{Q}^U \geq Q_j > Q_{j^*}$ (Region B in Figure 4). Once again, let a set of unreliable households of measure u_{jn} move out of pool j^* and contribute Q_j to pool j, for $j > j_d$ as defined above. Defining $\tilde{x}^U$ as the point where the Q_j-quantity line intersects $I^U(\bar{x})$, $\varepsilon_j(n) > 0$ can be chosen so to satisfy

$$\frac{u_{jn}Q_j p_U + \varepsilon_j(n)d}{u_{jn}Q_j + \varepsilon_j(n)} = K_j(n) = K(\tilde{x}^U)$$

for all $d \in (K(\hat{x}^U), \bar{p}]$. As for pools j contained between j_d and j^*, their only contribution will be given by the external agent: since by construction $K(\tilde{x}_j^U) > d$, no unreliable households will have incentives to deviate from the pooling allocation and contribute to these pools.

Next, take j for which $Q_{j^*} > Q_j > \hat{Q}^R$ (Region C in Figure 4). Let a set of reliable households of measure r_{jn} move out of pool j^* and contribute Q_j to pool j. Defining $\tilde{x}^R$ as the point where the Q_j-quantity line intersects $I^R(\bar{x})$, it is possible to choose $\varepsilon_j(n) > 0$ to satisfy

$$\frac{r_{jn}Q_j p_U + \varepsilon_j(n)d}{r_{jn}Q_j + \varepsilon_j(n)} = K(\tilde{x}^R)$$

for all $d \in [p_U, \bar{p}]$.

Finally, take j for which $Q_j \leq \hat{Q}^R < Q_{j^*}$ (Region D in Figure 4). In this case, the only one contributing to the pools is the external agent, whose delivery is $d \in [p_U, \bar{p}]$. Clearly, reliable households will not want to contribute anything to these pools.

Note that in all these cases then, household optimality holds and $K_j(n)$ is justified by using a reasonably pessimistic tremble. Then, it is confirmed that a perturbation satisfying these requirements supports the pooling allocation as a refined equilibrium.

It should be clear by now why it is necessary that unreliable households prefer the pooling allocation to the Rothschild-Stiglitz separating one. If that is not the case, then $K(\hat{x}^U)$ as defined above must be smaller than p_U, meaning that some inactive pools would have to offer an anticipated return below p_U if the pooling allocation is to be supported as a refined equilibrium. This is clearly impossible, since p_U is the lower bound of both households' and the external agent's contributions when a reasonably pessimistic tremble is used. ■