BEBR

#20-155

Coalition Proof Equilibrium In An
Adverse Selection Insurance Economy

Charles Kahn
Department of Economics
University of Illinois

Dilip Mookherjee
Statistical Institute
New Delhi

COALITION PROOF EQUILIBRIUM IN AN ADVERSE SELECTION INSURANCE ECONOMY

Charles Kahn and Dilip Mookherjee[*]

July 1991

ABSTRACT

We extend the notion of Coalition Proof Nash Equilibrium to a class of
matching games with private information. This solution concept is applied to
an adverse selection insurance economy, and is shown to yield a unique
allocation: the separating allocation without cross-subsidy. The relation to
alternative approaches to modelling contracting games with private
information, such as the Incentive Compatible Core, is discussed.

I. INTRODUCTION

This paper extends the notion of Coalition Proof Equilibrium (CPE) to a class of games where players possess payoff-relevant private information. Our extension can either be regarded as a generalization from the complete information context of the original recursive formulation of Berhneim, Peleg, Whinston (1987) or an extension of the stable set characterization used by Greenberg (1989) and Kahn-Mookherjee (1990). In the imperfect information context, it is necessary to redefine the notion of "blocking" of one coalitional agreement by another. The main difference is that blocking deviations are required to be "credible", in the sense that members of a coalition must signal their private information credibly to one another.

We apply the solution concept to a contracting game in the simple adverse selection insurance economy analyzed by Rothschild-Stiglitz (1976), Wilson (1977) and many others since. Our approach stands in contrast to purely noncooperative or cooperative approaches pursued in previous literature. In purely noncooperative approaches outcomes are very sensitive to the precise specification of the sequence of moves in the game (see, for instance, Hellwig (1987) and Kreps (1990)). They have been criticized by some authors (e.g., Boyd, Prescott and Smith (1988)) for imposing a given "organizational structure", interpreted as a particular extensive form that tend to rule out certain forms of communication and coordination among private agents. The latter criticism is particularly pertinent to the welfare implications that one may derive from such models.

These observations have motivated the study of cooperative approaches, where coalitions of economic agents may be interpreted as forms of

non-market organizations.[1] Nevertheless, the presence of private information imposes certain incentive compatibility restrictions on allocations that are feasible for any given coalition. These restrictions presume a noncooperative element that is also essential, so that the resulting theory needs to blend cooperative and noncooperative elements judiciously.

Our theory may be viewed as belonging to this genre. There are two main differences from previous literature. Most of the existing literature extends the notion of the Core to an Incentive Compatible (IC) Core,[2] where feasible allocations for a coalition must satisfy *individual incentive compatibility* restrictions, requiring that a proposed allocation should be immune to further deviations by singleton subcoalitions. Consequently equilibrium allocations have to satisfy a far more stringent requirement than alternative blocking proposals: the former have to be immune to deviations by all subcoalitions, whereas the latter must be immune to deviations only by singleton subcoalitions. In contrast, the approach we develop requires any blocking proposal to satisfy *coalition incentive compatibility* constraints, not just individual ones. These constraints apply in a consistent fashion to equilibrium allocations as well as proposed blocking deviations.

The second major difference concerns the notion of blocking. This implies different underlying assumptions regarding the process by which coalitions form, and the nature of monitoring and trading within any coalition. We assume that it is possible for a subcoalition to deviate away from an ongoing coalitional arrangement, without other members being aware of this deviation. Consequently, remaining members of the coalition cannot reorganize their trades or contracts in light of the deviation. In other words, contracting is decentralized in the sense that agreements between members of any given (sub)coalition cannot be monitored by those not belonging to it. In contrast,

the IC Core approach assumes that any deviating proposal is publicly observable, and that those not included in the deviating coalition can thereby reorganize their actions in light of the deviation.

As will become apparent in the application to the insurance economy, this assumption has important implications for the nature of equilibrium outcomes. If deviations are "private," then this destroys the viability of cross-subsidized contracts (since new contracts can be designed to attract away the profitable customers). If deviations are not private, then the original firm offering cross-subsidized contracts would be able to react in a way that may destroy the viability of the deviating contracts. Consequently, cross-subsidized contracts are viable when deviations are "public," not otherwise. The dependence of equilibrium outcomes on the ability of incumbent firms to react to new contracts has already been noted in the literature on noncooperative models of the insurance market. For instance, the different notions used by Rothschild-Stiglitz (Nash Equilibrium), Wilson (Anticipatory Equilibrium) and Riley (1979) (Reactive Equilibrium) correspond to different implicit assumptions regarding the ability of incumbent firms to react, or anticipate reactions of others. In many ways, the solution we obtain resembles the Riley notion of Reactive Equilibrium.[3]

It follows that "cooperative" approaches to modelling private information economies are also affected by the difficulty previously ascribed to noncooperative formulations: equilibrium outcomes can depend rather sensitively on the exact process by which contract proposals are made and reacted to. Nonetheless there does appear to be an advantage in the cooperative formulations that the dependency is tied to underlying notions of observability of actions rather than to details of timing in an artificially imposed set of rules. In our context, the welfare properties of equilibria

may depend on these aspects of the institutional process in an intrinsic manner, rather than on ad hoc restrictions on the set of coalitions that are allowed to form.[4]

Section II provides the extension of CPE to a general class of games with private information. Section III applies this solution concept to the simple insurance economy. Section IV concludes by discussing some implications of the analysis for the possible role of governments in enhancing welfare.

II. THE SOLUTION CONCEPT.

Let P denote the set of players. Players differ from one another in a number of characteristics, some of which are publicly observable, while others are known privately to the player concerned. Let T denote the set of *types* of individuals, assumed finite.[5] Types are partitioned into *categories*; two types which belong to the same category are publicly indistinguishable. In other words, everyone in the economy can identify a person's category, but only he himself knows his own type. Let N denote the set of categories. Let T_i denote the set of all possible types in category i. Any player in category i has available to him the set of actions A_i.

The underlying game is played by a set of players which is constructed by selecting one member from each category The selection of the set of players is made by a random matching process. Suppose that the matching process has brought together a group of players from different categories, where the category i player happens to be of type t_i. Let $t = (t_1, t_2, ..)$ denote the resulting vector of types in this player group. If in this particular play of the game the set of chosen actions is $a = (a_1, a_2, ...)$, then the player of category i obtains a von Neumann Morgenstern payoff $U(a|t)$.

5

Notice that the payoff depends not only on the actions of other players but also of their types. The crucial limitation to this selection device is that a player's payoff is independent of the actions chosen by other players in the same category.[6]

For simplicity, we assume throughout the analysis that all players of a given type behave identically, both in choice of action and in decisions about joining a coalition. Effectively, this simplification is a restriction to pure strategies. With this assumption, we can henceforth think of a player as a representative member of any given type. A *strategy for category i* is therefore a function $s_i : T_i \rightarrow A_i$.

There are two interpretations of the matching process. Either it is a one-time selection, where the underlying game is played just once by the selected group, or it is a repeated selection, in which case each player is assumed to make the same play in all encounters. Under the first interpretation of the matching process, let $q_i(t_{-i}|t_i)$ represent the probability assessment made by a player of type t_i that he will be grouped with a set of players from other categories whose types happen to be represented by the vector t_{-i}. (Under the second interpretation, it represents the relative frequency that type t_i is matched with a set with type vector t_{-i} when groups are selected to play the game repeatedly.) We define the ex ante utility of a player of type t_i as a function of the vector s_{-i} of strategies selected by players in other categories, and the action a_i chosen by this type:

$$W_i(s_{-i}, a_i | t_i) = \sum_{t_{-i} \in T_-} q_i(t_{-i}|t_i) \ U_i(s_{-i}(t_{-i}), a_i | t_{-i}, t_i) \qquad (2.1)$$

In effect we have reformulated the game of incomplete information as a

6

Bayesian game. This enables us to define a strategy vector $s = (s_1, s_2, \ldots)$ to be a *(Bayesian) Nash Equilibrium* if for all $i \in N$ and all $t_i \in T_i$:

$$s_i(t_i) \in \underset{a_i \in A_i}{\text{argmax}} \quad W_i(s_{-i}, a_i | t_i) \tag{2.2}$$

In what follows, we will use the repeated encounter interpretation of the matching technology, so that there are a large number of players of the different types and the matching represents a repetition of independent draws. It is for this particular interpretation that the following description of the preplay communication process is most apt. It is also the natural interpretation for application to the insurance market game in the following section, where different insurance firms and the set of insurance customers constitute the different categories of players, and there are a large number of players in the customer category (divided into different risk types).

Before play begins, different coalitions may form in order to coordinate on their strategy choices. A *coalition* is a subset of the set of all players P, with the restriction that if a player of a given type is a member of that coalition, then so are all other players of the same type. This restriction ensures that a coalition is simply a subset of T, i.e. is a collection of types, typically from different categories. Of course, it is not possible for any member of a coalition to verify the types of other members. We shall describe below how members of a coalition can infer the types of other members: specifically, we shall impose a credibility constraint on any agreement that a coalition may seek to enter into, where members signal their types to each other by their willingness to be party to the agreement.

An *agreement* (or *contract*) is a combination of a coalition C and a strategy vector $s = (s_1, s_2, \ldots)$. It may be interpreted as an agreement among

members of C to coordinate on the choice of actions as stipulated by s, given their common expectation that non-members are choosing actions as stipulated by s.

An agreement (C,s) is said to *undermine* the agreement (D,v) (denoted $(C,s) \gg (D,v)$) if:

(i) $C \subset D$

(ii) $s_i(t_i) = v_i(t_i)$ for all $t_i \notin C$

(iii) All members of C have higher ex ante payoffs at s than at v:[7]

$$W_i(s|t_i) > W_i(v|t_i) \text{ for all } t_i \in C.$$

(iv) The agreement (C,s) is a *credible deviation* from the strategy vector v, i.e. for any $t_i \in C$, any other type $\hat{t}_i \in T_i$ does not belong to C if and only if

$$\sup_{a_i \in A_i} W_i(v_{-i}, a_i | \hat{t}_i) \geq \sup_{a_i \in A_i} W_i(s_{-i}, a_i | \hat{t}_i) \qquad (2.3)$$

The first condition requires the deviating coalition to be a proper subset of the original coalition.[8] In particular this implies that an agreement cannot be undermined by a group of individuals some of whom are not party to this agreement. This is consistent with the spirit of all agreements (not involving the grand coalition P) constituting private deviations, where those outside the coalition are not aware of the deviation taking place. For instance, if D is a proper subset of P, then those not in D would not be aware that members of D were intending to play as in v. In such cases they could not possibly be involved in engineering a further deviation that would undermine (D,v).

The second and third conditions are self evident: the second requires that only members of the deviating coalition change their action choices,

while they expect nonmembers to continue to choose actions as in the original agreement. (iii) requires the deviants to be better off in expectation.

The fourth condition deserves some explanation. Note that members of a given coalition cannot verify each others' types. Nevertheless, the profitability of the deviation for any member will in general depend on the types of his partners. Consequently, members will be concerned to assure themselves that their partners are really of the types that they claim to be.[9] Condition (iv) is a natural condition that ensures that the types of different members are credibly signalled to one another. It says that if there is a player from a certain category i in the coalition who claims to be of type t_i, then this player can convince others that he is not of a type $\hat{t}_i$ that is not represented in the coalition. For if he were of this alternate type $\hat{t}_i$, then the highest expected utility that he could possibly get following the deviation is given by the right hand side of (2.3). This does not exceed what such a type would be getting in the status quo (the left hand side of (2.3)). In other words, such a type could not conceivably benefit from the deviation, and therefore could not provide support to it. In contrast, all types ostensibly included in the deviation do strictly benefit from it (as ensured by condition (iii)).[10] Our criterion therefore bears a close relation to the Intuitive Criterion of Cho and Kreps (1987) for refining sequential equilibria of noncooperative extensive form games.

At issue in any game is the question of which agreements will be viable. For an agreement (C,s) to be viable, it should not be undermined by a deviating agreement involving a subcoalition of C. Invoking the principle of *consistency*, we require deviating agreements themselves to be inviolate against deviations by subgroups of the initial subcoalition. This leads to the following recursive definition of a *self-enforcing agreement*:

If the set C contains a single type, then for any strategy vector s, the agreement (C,s) is *self-enforcing*.

If the set C contains more than one type, the agreement (C,s) is *self-enforcing* if for no proper subset D of C is it the case that there is a self-enforcing agreement (D,v) which undermines (C,s).

A self-enforcing agreement is therefore the natural formulation of *coalition incentive compatibility*: If (C,s) is self-enforcing, then there does not exist a subcoalition that can profit from a self-enforcing deviation. The following, more concise definition makes use of the von-Neumann Morgenstern concept of an *abstract stable set*:

A stable set of agreements is a set G such that

$$(C,s) \in G \Longleftrightarrow \text{there does not exist } (D,v) \in G \text{ such that } (D,v) \gg (C,s).$$

In the example we investigate in the next section a unique stable set of agreements exists, namely the set of self-enforcing agreements.[11]

Either definition of self enforcing agreements can be used to generate our equilibrium concept:

A strategy vector s is said to be a *Coalition Proof Equilibrium with Private Information* (CPEPI) if:

(i) the agreement (P,s) is self enforcing, and

(ii) there does not exist any other self enforcing agreement (P,v) which Pareto-dominates (P,s).

10

It can be shown that this definition reduces to BPW's Coalition Proof Equilibrium in the absence of private information provided that the set of players is finite and the set of actions for each player is finite. It can also be verified that a CPEPI is always a Bayesian Nash equilibrium.[12]

III APPLICATION TO THE INSURANCE ECONOMY

We first describe the economy, and later explain how we cast it in terms of the game theoretic model of the previous section. There are two kinds of economic agents: insurance customers and insurance firms. Each customer belongs to one of two risk classes: high (H) or low (L). Each customer knows his own risk class, but no firm can identify the risk class of any specific customer. There are a large number of customers, and a fraction $\mu_i > 0$ of the customer population is commonly known to be of risk class i (i = H,L), where $\mu_H + \mu_L = 1$.

For each customer, there are two states of nature: accident (A) and no accident (NA). In the accident state the customer suffers a loss of endowment. There is a single physical commodity, the endowment loss is d units of this commodity. There are therefore two contingent commodities for every consumer, corresponding to delivery in the two states of nature, and the commodity space can be identified with $\mathbb{R}^2$. Letting the first commodity correspond to delivery in the no accident state, the endowment of every customer is (0,-d), where d > 0.

All customers share the common von Neumann Morgenstern utility function $U:\mathbb{R} \to \mathbb{R}$ which is continuously differentiable and strictly concave (so every

customer is risk averse). Use p_i to denote the probability of the accident state for risk class i, where $1 > p_H > p_L > 0$. So either state could arise for either type of customers, but the likelihood of the accident state is higher for the high risk types. If the type i customer enters into a net trade of (x_0, x_1) his expected utility will therefore be

$$U_i(x_0, x_1) \equiv p_i U(x_1 - d) + (1 - p_i) U(x_0) \qquad (3.1)$$

There are a number of risk neutral insurance firms, denoted f = 1,2,3,..F.[13] If firm f provides a net trade vector (x_0^f, x_1^f) to a customer of risk type i, it will obtain an expected profit of

$$\pi_i(x_0^f, x_1^f) \equiv -p_i x_1^f - (1 - p_i) x_0^f \qquad (3.2)$$

======================= INSERT FIGURE 1 =======================

Figure 1 represents indifference curves of the two types of customers, as well as isoprofit curves for firms. The origin depicts the initial endowment point; the line ABC is the locus of points of full insurance. Line OCD is the locus of net trade vectors that would generate zero profits for supplying firms when offered to high risk customers, and OAF is the corresponding line for low risk customers. Indifference curves for high risk customers are flatter than for low risk customers, and are tangent to corresponding isoprofit lines along the line of full insurance.

The complete information Walrasian outcome is for a customer of each risk type to obtain full insurance at actuarially fair odds, i.e. so H types would receive C and L-types A. In the absence of information about risk type, this

allocation would fail to be incentive compatible: high risk customers would masquerade as low risk in order to avail of the trade A rather than C.

A *separating allocation without cross-subsidy* is a trade allocation in which the high risk customers obtain full insurance at actuarially fair odds (i.e. the point C in Figure 1), and low risk customers obtain incomplete insurance, also at actuarially fair odds, at the point (D in Figure 1) where the high risk customers are indifferent between the two trade vectors. These trade vectors are hereafter denoted x_H and x_L respectively.

We now describe the trading game for this economy, in terms of the general class of games considered in the previous section. The set of categories is $N = \{0,1,2,3,..F\}$. Players of category 0 are insurance customers, while category $i \geq 1$ corresponds to firm i. Thus, there are a large number of players in category 0, but a single player in category $i \geq 1$. Players in category 0 belong to two risk types, so $T_0 = \{H,L\}$. On the other hand, there is perfect information regarding firms, so for $i \geq 1$, T_i contains a single element. Moreover, the matching process is represented by

$$q_i(t_{-i}|t_i) = \left\{ \begin{array}{ll} \mu_H & \text{if } t_0 = H \\ \\ \mu_L & \text{if } t_0 = L \end{array} \right.$$

for any $i \geq 1$.

Customers formulate trade requests, while firms make trade offers. Trade occurs when request and offer match. Player $f \geq 1$ has the action set

$$A_f = \{ S^f \mid S^f \subseteq \mathbb{R}^2 \}$$

so firm f selects a set S^f of trade offers.

13

The consumer chooses his action $(f*, r)$ from the set of trade requests:

$$A_0 = \{1, 2, \ldots F\} \times \mathbb{R}^2$$

in other words, the customer selects a firm f^* and makes a request r.[14] A trade takes place only if the consumer requests it and the firm offers it:

$$x^f = \begin{cases} r & \text{if} \quad f = f^* \text{ and } r \in S^f \\ \\ 0 & \text{otherwise} \end{cases}$$

Consequently the payoff of type j of player 0 is:

$$U_0(a_0, a_1, a_2, \ldots \mid t_0 = j) = p_j U(\sum_f x_1^f - d) + (1 - p_j) U(\sum_f x_0^f)$$

and for player $f \geq 1$ is:

$$U_f(a_0, a_1, a_2, \ldots \mid t_0 = i) = -p_j x_1^f - (1 - p_j) x_0^f$$

We use $V_j(x)$, $j = H, L$ to denote the expected utility obtained by a the customer from the net trade vector x.

PROPOSITION   *There is a unique CPEPI outcome which is the separating allocation without cross subsidy.*

Proof: It suffices to establish that the strategy vector s with the property that (P,s) is self enforcing, are strategy vectors which yield separating

allocations without cross-subsidy.

STEP 1: <u>There exists a self enforcing agreement (P,s) resulting in the separating allocation without cross subsidy.</u>

Consider the following strategies: H types of the customer request x_H from firm 1, L types request x_L from firm 1, firm 1 offers $S^1 = \{x_H, x_L\}$, while firms $f \geq 2$ offer nothing. We claim that this is self-enforcing. It is evident that every single player is playing a best response, so (P,s) cannot be undermined by any singleton coalition. Moreover, the only coalitions that can conceivably undermine (P,s) must include some firms, and one or both types of the customer.

Suppose there exists a coalition involving some firm and only the H type of the customer, which can deviate to a self-enforcing agreement which undermines (P,s). Then H types must obtain a higher expected utility than at x_H, and must therefore be able to secure a trade vector which generates losses for whichever firm offers it. Since L types do not form part of the coalition, they must continue to request the trade x_L from Firm 1. So the new trade obtained by the H types following the deviation is obtained only by the H types. The deviating agreement cannot be self-enforcing, then, because it can be undermined in turn by a further deviation where the firm offering the new trade to the H types drops this offer and avoids the resulting losses.

Suppose (P,s) is undermined by a coalition involving some firm(s) and only L types of the customer. In order for this deviation to be credible, it must be the case that the H types do not prefer the new trade obtained by the L types. This implies that the new trade obtained by the L types must lie in the shaded region of Figure 1. Such a trade would generate losses for the firms selling it, and (using an argument analogous to that in the previous

15

paragraph) the deviating agreement cannot therefore be a self-enforcing.

Suppose (P,s) is undermined by an agreement involving a coalition involving both types of customers, and some firms. Then both types of customers must be better off in the agreement. Consequently the trade obtained by the H type must lose money for whichever firm offers it. If this trade differs from the trade obtained by the L type, then it is clear that the agreement cannot be self-enforcing since the firm offering the trade to the H type will be better off dropping the offer. Hence the deviation must involve a pooled trade, at a point like Q in Figure 2. Moreover, the deviation must involve a single firm (since all parties to the deviation must be strictly better off, and all types of the customer trade with a single firm).

-------------------- FIGURE 2 ------------------------

We claim that the deviating agreement is not self-enforcing. It is evident that the trade Q loses money on the H types, and earns profits on the L types. Moreover, the marginal rates of substitution of the two types necessarily differ at Q. We claim that the deviating agreement is undermined by the following self-enforcing agreement between the same firm and the L types: the firm withdraws its offer of Q, and offers instead the trade R (close to Q), where $V_L(R) > V_L(Q)$ and $V_H(R) < V_H(Q)$. At the same time the L types request the trade R from this firm. It is evident that the deviating agreement is credible, and undermines the original deviation. Moreover, both the concerned firm and the L types are playing best responses, so the deviating agreement is self-enforcing. Hence the original deviation cannot be, concluding the proof of Step 1.

16

STEP 2: There cannot be a self enforcing agreement (P,v) where H types get lower expected utility than at the separating allocation without cross-subsidy.

Suppose this is false. Then there exists a point of full insurance (such as G in Figure 2) close to $C = x_H$, which earns positive expected profit for both types of customers, and which generates higher expected utility for the H types. If the L types are at least as well off in v as at the trade G, a coalition of an inactive firm (i.e. a firm that does not trade with either type of customer in v) and the H types could form an agreement to deviate to a trade of G. It is evident that this agreement is credible, self-enforcing and undermines (P,v).

So consider the case that L types are better off at G than they are in v. Then a coalition containing an inactive firm, and both types of customers, could agree to deviate by trading G (i.e. the firm would offer G and both types would request G from the firm). We claim that the deviating agreement is self-enforcing. No singleton subcoalition could profitably deviate. Moreover, a coalition of the firm and the H type alone would not be able to find a self-enforcing agreement to undermine the deviation. Since the deviation does not involve the L type, it must be the case that the L type continues to request G from the firm. Since the sale of G to the L type is profitable for the firm, it must be the case that a self enforcing deviation will involve the firm continuing to offer G. Hence the firm must offer some additional trade, M, say which is requested by the H type, and this must generate higher expected profit for the firm than trading G with the H type. Moreover, the H type must also be better off at M than at G. No such trade M can exist, since G is a point of full insurance.

A symmetric argument establishes that a subcoalition of the firm and the

L types cannot also find a self-enforcing deviation to undermine the original deviation. So the original deviating agreement is self-enforcing, establishing that (P,v) is not self enforcing.

STEP 3: <u>There cannot be a self enforcing agreement (P,v) resulting in a separating allocation where the H types obtain higher expected utility than at $\underline{x}_H$.</u>

If this were false, the trade obtained by the H type would earn losses for the firm trading with it (since by assumption the L types obtain a different trade), and this firm would be better off dropping the offer of this trade.

STEP 4: <u>There cannot be a self enforcing agreement (P,v) resulting in a pooled allocation.</u>

Consider first the possibility that (P,v) results in both types obtaining the trade $C = x_H$. Then a coalition of an inactive firm and the L types could deviate to an agreement to trade a point such as J in Figure 2. This is evidently a credible deviation which is self enforcing and undermines (P,v).

Using the result of Step 2, a pooling allocation resulting from a self enforcing (P,v) must therefore lose money on the H types. The firm offering the trade must break even, and must therefore be earning positive profit from the L types. (P,v) would then be undermined by a self-enforcing deviation involving the firm and the L type (for example as towards the end of the reasoning in Step 1).

STEP 5: <u>Every agreement (P,v) that is self enforcing must result in the non-cross-subsidized separating allocation.</u>

By Step 4, a self enforcing agreement (P,v) must result in a separating

18

allocation. By Steps 2 and 3, and the fact that the trade offered to the H type must break even, all such agreements must result in the H type obtaining a trade of $C = x_H$. The trade obtained by the L type must also earn nonnegative expected profit, and allow the L type at least as much expected utility as C. Moreover, the H type must also not envy the trade obtained by the L type. Now note that the trade $D = x_L$ maximizes the expected utility of the L type subject to the constraint that the resulting allocation (i.e. combined with C for the H type) is incentive compatible, and that the trade earns nonnegative expected profit. If the L type receives a trade with a lower expected utility, he can form a coalition with an inactive firm and deviate to an agreement to trade a point in the neighborhood of D which earns positive expected profit, generates a higher utility for the L type, and is not envied by the H type. This agreement would be good and would undermine (P,v). Hence every self enforcing (P,v) would result in the allocation (x_H, x_L). Q.E.D.

Self enforcing agreements do not permit the use of cross-subsidized contracts, for the reason that such contracts are vulnerable to "creamskimming," i.e. an inactive firm can offer a contract which attracts away only the low risk customers. This is quite similar to the way a pooling allocation is not sustainable in the Rothschild-Stiglitz analysis. In a cross subsidized separating allocation has the additional problem that the firm offering the trade to the high risk customers will be better off withdrawing this trade.

In contrast, the IC Core approach of Boyd, Prescott and Smith (1988) and Marimon (1988) allow cross subsidized allocations to be viable. This is the crucial distinction between the two approaches. For instance, Marimon requires every deviating proposal to respect "property rights" of the complementary

19

coalition, which has the effect of ruling out creamskimming deviations (since the original firm left selling to the high risk customers alone would suffer losses as a result of the deviation). Boyd, Prescott and Smith require the deviating proposal to satisfy the property that members of the complementary coalition should not prefer to join the deviating coalition: this also has the effect of prohibiting creamskimming.[15] The IC Core approach thus pertains to a context where all deviating proposals are publicly observable, in contrast to our approach which presumes that deviations occur privately. The distinction may be thought of as corresponding to two different trading environments, with the IC Core relevant to a centralized framework with a richer structure of public information. Our approach pertains instead to a decentralized trading arrangement, the hallmark of which is that contracts between a set of individuals are unobservable to others.

Despite the possibility of creamskimming, our analysis does not suffer from problems of nonexistence akin to the Rothschild Stiglitz analysis. This is where the consistency of the solution concept is essential. While any pooled contract can be undermined by a separating contract, a separating contract cannot be undermined by a pooled contract (since the latter is itself non-viable). Since an undermining agreement is required to be a good agreement, players are credited with a certain degree of foresight. It is this feature of our solution concept that makes it similar to the concept of Reactive Equilibrium introduced by Riley (1979).[16]

It is well known that the separating allocation without cross subsidy is constrained inefficient if the proportion of high risk customers in the population is low (see Rothschild-Stiglitz (1976, pp 644-5) or Crocker and Snow (1985)). For instance, in cases where there is no Nash Equilibrium in the Rothschild-Stiglitz model, this allocation is Pareto dominated by any pooling contract that breaks it. It pays low risk customers to offer to subsidize high risk customers if that relaxes the incentive constraint for the latter sufficiently to relax the constraint on the quantum of insurance available to the low risk customers. If the proportion of high risk customers is small, the per capita cost of a given subsidy for any given low risk customer is small, and a Pareto improvement results.

As demonstrated in the previous section, a market outcome is not consistent with cross subsidization, since such arrangements are vulnerable to creamskimming. Inefficiency thus results from the decentralized character of the contracting process.[17] This may be alternatively interpreted as a problem of insufficient commitment abilities on the part of market agents. This problem appears to be peculiar to adverse selection environments: with a similar model of contracting in a moral hazard insurance environment. Coalition Proof Equilibrium outcomes correspond exactly to the set of constrained efficient outcomes (see Kahn and Mookherjee (1990b)).

What specific policies may the government adopt in order to effect welfare improvements? A natural candidate is the public provision of cross-subsidized or pooled insurance. It may appear that such schemes are also vulnerable to creamskimming: private firms may have an incentive to make profits by weaning away the low risk customers from the public pool. This

problem, however, does not arise if public insurance is provided universally and compulsorily, i.e. low risk customers do not have the freedom to opt out of the public insurance program (by paying lower taxes). In particular, note that it is feasible for the government to implement the pooled insurance program (point B in Figure 1) which breaks even for the population as a whole, and provides each customer with full insurance (assuming the absence of moral hazard). This has the effect of shifting the initial endowment of all customers to a point where all gains from trade between customers and private firms are exhausted. Such an outcome is always constrained efficient: indeed it is efficient in a first best environment, and it Pareto dominates the CPEPI outcome whenever the latter is constrained inefficient.

It is the assumed superior ability of the government to commit that accounts for its potential welfare role in the previous analysis. A private firm would not be able to credibly replicate the policy of providing the pooled full insurance policy that breaks even because of its inabilility to commit to continue providing this policy: it will have a temptation to drop its high risk customers to increase its profits. The government's superior ability to commit may stem from the different nature of its objectives: concern for welfare or reelection prospects may cause it to be more sensitive to the comprehensiveness of public insurance programs than to the profitability of such programs. A more explicit analysis of commitment possibilities, however, is a task for future research.

REFERENCES

Allen, B. (1991), "Market Games with Asymmetric Information: The Value," Presented at the Midwest Mathematical Economics Meetings, Northwestern University, May 31, 1991.

Berliant, M. (1990) "On Income Taxation and the Core," May 1990, working paper, University of Rochester.

Bernheim, D., B. Peleg and M. Whinston (1987), "Coalition Proof Nash Equilibria I. Concepts," *Journal of Economic Theory*, 42, 1-12.

Boyd, J., E. Prescott and B. Smith (1988), "Organizations in Economic Analysis," *Canadian Journal of Economics*.

Boyd, J. and E. Prescott (1986), "Financial Intermediary Coalitions," *Journal of Economic Theory*.

Chakravorti, B. and C. Kahn (1990), "Universal Coalition Proof Equilibrium," mimeo, University of Illinois, Urbana Champaign.

Cho, I.K. and D. Kreps (1987), "Signalling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102, 179-221.

Crocker, K.J. and A. Snow (1985), "The Efficiency of Competitive Equilibria in Insurance Markets with Asymmetric Information," *Journal of Public Economics*, 26, 207-219.

Greenberg, J. (1989), "Deriving Strong and Coalition Proof Nash Equlibria from an Abstract System," *Journal of Economic Theory*, 195-202.

Hellwig, M. (1987), "Some Recent Developments in the Theory of Competition in Markets with Adverse Selection," mimeo, University of Bonn.

Jaynes, G.D. (1978), "Equilibria in Monopolistically Competitive Insurance Markets," *Journal of Economic Theory*, 19, 394-422.

Kahn, C. and D. Mookherjee (1990a), "The Good, the Bad and the Ugly:

Coalition Proof Equilibria in Infinite Games," forthcoming, *Games and Economic Behavior*.

_____ (1990b), "Decentralized Exchange and Efficiency under Moral Hazard," mimeo, University of Illinois, Urbana.

_____ (1991), "Optimal Incentives for Agents with Side Contracting Opportunities," mimeo, University of Illinois, Urbana.

Kahn, J., (1987) "Endogenous Financial Structure in an Economy with Private Information," University of Rochester working paper.

Krasa, S. and N. Yannelis, (1991) "Value and Private Information," University of Illinois working paper.

Kreps, D. (1990), *A Course in Microeconomic Theory*, Princeton Univ.Press.

Lacker, J. M., and J. A. Weinberg, (1990) "A 'Coalition Proof' Equilibrium for a Private Information Credit Economy," Working Paper 90-8, Federal Reserve Bank of Richmond.

Marimon, R. (1988), "The Core of Private Information Economies," mimeo, University of Minnesota.

Miyazaki, H. (1977), "The Rat Race and Internal Labor Markets," *Bell Journal of Economics*, 8 (1977), 394-418.

Pauly, M. (1974), "Overprovision and Public Provision of Insurance," *Quarterly Journal of Economics*, 93, 541-562.

Riley, J. (1979), "Informational Equilibrium", *Econometrica*, 47, 331-59.

Rothschild, M. and J. Stiglitz (1976), "Equilibrium in Competitive Insurance Markets," *Quarterly Journal of Economics*, 90, 629-649.

Wilson, C. (1977), "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory*, 16, 167-207.

Yannelis, N. (1991), "The Core of an Economy with Differential Information," *Economic Theory*, 1.

[1]An early interpretation of equilibrium arrangements with private information as nonmarket organizations is provided by Miyazaki (1977). Boyd, Prescott and Smith (1988) develop this theme in particular. Other cooperative analyses of private information economies include Boyd and Prescott (1986), Kahn (1987), Marimon (1988), and Berliant (1990).

[2]See however, Yannelis (1991), Krasa-Yannelis (1991) and Allen (1991) for extensions of the value to private information contexts. Chakravorti-Kahn (1990) consider an extension related to the bargaining set. Lacker and Weinberg (1990) consider an IC-Core extension which imposes the same consistency requirements on blocking and equilibrium allocations.

[3]To the extent that we assume that nondeviants cannot react to any given deviation, this may appear surprising. One would expect that we would continue to be plagued by the nonexistence of equilibrium as did Rothschild and Stiglitz. However, we do obtain existence of equilibrium by insisting on *consistency* of the solution concept, i.e. that deviating proposals be subject to similar restrictions as are proposed equilibria. Since cross-subsidized contracts are not viable themselves, they cannot be used to block non-cross-subsidized separating allocations. In this manner the Rothschild-Stiglitz nonexistence problem disappears. Morever, agents proposing a deviating contract must 'look ahead' and check whether this contract can be rendered unprofitable by a further deviation, before considering it seriously. For this reason, we end up with a solution concept that resembles the Reactive equilibrium notion.

[4]Recall that constrained efficient allocations often involve cross-subsidization. The solutions generated by our model may therefore fail to be constrained efficient.

[5]We describe the model with a finite set of player categories and types for simplicity. In some contexts an infinity of player categories may be necessary, as in the free-entry application of Kahn-Mookherjee (1990b). In this case the stable sets approach can be used to generalize the equilibrium concept; see Kahn-Mookherjee (1990a).

[6]Even this limitation is not absolute: for applications in which individuals who are of identical categories have effects on one another, the game can be modified by replicating the set of categories.

[7]We abuse notation by using $W_i(s|t_i)$ to denote the expected payoff of type t_i at the strategy vector s. In terms of expression (2.1) for expected payoffs, we have $W_i(s|t_i) = W_i(s_{-i}, s_i(t_i)|t_i)$.

[8]The standard formulation of Coalition Proof Equilibrium (BPW (1987), Greenberg (1989), Kahn-Mookherjee (1990a)) specifies weak rather than strict inclusion. Kahn-Mookherjee (1990a) show the advantages of this alternate formulation in general; in the context examined here the results are equivalent but it simplifies the presentation.

[9]Strictly speaking, they will need to assure themselves that the actual coalition does not contain any type that is not supposed to be part of it.

[10] It is important to remember that excluded types are actually unaware of the deviation. Condition (2.3) deals with a counterfactual essential to make the deviation credible: included types convince each other of their identities by demonstrating that they couldn't possibly benefit from the new agreement if they were actually of a different, unrepresented, type.

[11] Kahn and Mookherjee (1990a) show that the stable set approach gives the better generalization of CPE in more complicated contexts. In general, a stable partition may not exist; in such situations the notion of a stable partition has to be extended to a semi-stable partition. That paper also deals with extensions when there are multiple stable sets.

[12] This follows from the fact that any deviation by a singleton coalition is always credible, since different types of the same player do not interact with one another, and a deviation by any one type therefore does not affect the payoffs of other types (so that (2.3) is always satisfied for any such deviation).

[13] F is at least 3.

[14] Thus customers are constrained to trade exclusively with a single firm. This is a significant restriction, because customers will typically have an incentive to trade with multiple firms. For a discussion of this issue, see Pauly (1974), Jaynes (1978) and Kahn and Mookherjee (1991). Nevertheless, most models of insurance markets assume that information about trades between any set of agents is publicly observable, so exclusive contracts are feasible. Morever, it is evident that any firm trading with any customer will always (weakly) prefer to be the exclusive provider. Thus if the technology exists to enforce exclusive trades, the insurers will make sure it is adopted.

[15]Boyd, Prescott and Smith do not model firms explicitly: the above comment pertains to the obvious extension of their model with firms as explicit players.

[16]One difference between the Reactive Equilibrium and our solution concept ought to be noted: in the former, a firm offering a contract foresees the reactions of *other* firms, while in the latter it can also foresee its *own* temptation to deviate away from the contract. In this respect, our notion also captures a criterion of renegotiation proofness.

[17]In practice, group insurance practices involving cross subsidization are commonly observed. Nevertheless, it is unclear to what extent this reflects purely voluntary behavior: for instance, employer-provided insurance is bundled with other attributes of an employment contract.

FIGURE 1

FIGURE 2