Equilibrium in Insurance Markets with Asymmetric Information and Adverse Selection

Jonathan A. K. Cave

Rand

HEALTH INSURANCE EXPERIMENT SERIES

The research reported herein was performed pursuant to Grant No. 061B80 from the U.S. Department of Health and Human Services, Washington, D.C.

84-4820
ISBN 0-8330-0554-5

Published by The Rand Corporation

R-3015-HHS

Equilibrium in Insurance Markets with Asymmetric Information and Adverse Selection

Jonathan A. K. Cave

April 1984

Prepared under a grant from the
U.S. Department of Health and Human Services

Rand
1700 MAIN STREET
P.O. BOX 2138
SANTA MONICA, CA 90406-2138

PREFACE

This report examines possible outcomes of greater competition in insurance markets. It applies directly to the situation where one insurer offers multiple options (e.g., high and low) to an employee group. Many observers argue that requiring employers to make multiple-option insurance available will improve efficiency of the market allocation. Although this report does not directly compare the no-options and multiple-options situations, that is the subject of one of the first analyses of the multiple-options case.

The report removes a rather restrictive assumption in the existing theory of insurance markets and is thereby able to explain observed phenomena in a manner hitherto not possible. Specifically, it describes the nature of insurance offerings in equilibrium if firms offer multiple policies; but it replaces the conventional assumption that each policy must earn nonnegative profits with the more realistic requirement that the portfolio of policies offered by the firm earn nonnegative profits in the aggregate. The equilibrium in general requires that certain policies (low option) subsidize others (high option). Theorems regarding the existence, optimality, and uniqueness of the subsidy equilibrium are presented, together with a simple characterization of the subsidy equilibrium and a comparison with existing equilibrium notions. Because the subsidy patterns, from low to high, that emerge under this formulation appear to characterize multiple-option insurance plans such as the Federal Employees Health Benefits Plan, this model may be more useful than conventional methods in the analysis of such plans.

This work was conducted under a grant from the Department of Health and Human Services as part of the Health Insurance Experiment. It is intended for policymakers interested in the welfare effects of competition under conditions of asymmetric information, and for economists and game theorists concerned with adverse selection problems and games of asymmetric information. The reader should have a thorough grounding in economics and some familiarity with Nash and perfect equilibria of noncooperative games.

SUMMARY

In analyzing the effect of adverse selection on insurance markets and other markets where one side of the market may have better information than another, it has been conventional to assume that contracts are "actuarially fair." In particular, when insurance policies are offered to different types of consumers, and the firms offering the insurance cannot observe the type or risk class of the consumer, it is generally assumed that *each* contract breaks even on average; there is no subsidization of one class by another. However, such subsidies appear to be an important feature of actual insurance offerings.

At the same time, the equilibria that firms reach under the nonsubsidy assumption have some disturbing features. First, if firms assume that their competitors will respond to the presence of a new policy by continuing to offer their current policies, even if they become unprofitable, there may well be no equilibrium. On the other hand, if firms adopt the more sophisticated point of view that new offerings will be greeted with discontinuance of unprofitable offerings, "pooling" equilibria may exist that lump risk classes together. However, under neither set of conjectures about the response of existing firms will the policy offerings necessarily be Pareto-optimal. They will also not deter entry by a sophisticated firm willing to sacrifice actuarial fairness for nonnegative overall profits.

We model a process of competition between firms who do not require that all of their policies break even. Instead they will offer portfolios of policies that earn nonnegative profits overall, when portfolios made unprofitable by entry have been withdrawn.

It is demonstrated that such behavior always results in an equilibrium for any distribution of the consumer population between a finite number of risk classes. Moreover, except for exceptional cases, this equilibrium is unique.

To compare these results with the no-subsidy results, it is easiest to think of a world with two classes of consumers: high-risk and low-risk. Conventionally, if the proportion of high-risk agents is high enough, there will be a separating equilibrium that offers the high-risk agents full fair coverage, and offers the low-risk agents the best fair coverage that will not attract high-risk agents. This is distinctly inferior to full fair insurance for the low-risk agents, so that there is a sense in which the low-risk agents suffer a negative externality. If the proportion of low-risk agents is sufficiently high, there will either be no equilibrium (with naive anticipations) or a pooling equilibrium (with sophisticated anticipations). Unlike the separating equilibrium the pooling equilibrium is never Pareto-optimal, but is sensitive to the proportions of high- and low-risk agents.

Dropping the no-subsidy requirement results in an equilibrium that is always Pareto-superior to the no-subsidy equilibrium, is Pareto-optimal, and is sensitive to the proportions of high- and low-risk agents. The low-risk agents subsidize the high-risk agents, and in exchange obtain better terms than they could otherwise achieve. In fact, the subsidy equilibrium is precisely the pair of policies that maximize the utility of the low-risk agents subject to the constraints of zero overall profits; full insurance of high-risk agents; and expected utility maximization by all agents (which separates types). The only exception to this is when the proportion of high-risk agents is so high that this would result in adverse subsidies. For proportions above this critical level (which strictly exceeds the critical proportion dividing separating from pooling equilibria), the subsidy and separating equilibria coincide.

The fact that this equilibrium seems to capture a feature of the real world that is absent from the no-subsidy model, and to enjoy the simplicity of a model that predicts the existence of a unique equilibrium as opposed to the multiplicity of "pooling" equilibria that can occur when there are more than two risk classes, suggests that it may be a more realistic model for analyzing the behavior of firms offering multiple-option insurance contracts.

CONTENTS

FIGURES

I. INTRODUCTION

This report examines some of the consequences of allowing firms that operate in markets characterized by asymmetric information between sellers and buyers to compete, using coordinated portfolios of offers that may involve subsidization.

Equilibria of markets with adverse selection have been investigated by many authors. The pioneering paper of Rothschild and Stiglitz (RS) described a model in which consumers know their own risk classes (accident probabilities), and firms do not. Their firms are constrained to offer single-option policies, and free entry forces each firm to make zero expected profits; hence, each policy breaks even or is "actuarially fair" for the group it serves. The model is solved by Nash equilibrium between the producers and expected utility maximization by the consumers. For the case of two types of consumers (high- and low-risk), the authors described a distinguished pair of contracts, independent of the population parameters, that would form the Nash equilibrium if one existed. However, they also showed that if the proportion of high-risk agents was sufficiently low, no Nash equilibria would exist.

The RS construction, including the definition of the "separating contracts" that form the Nash equilibrium when it exists, can easily be generalized to arbitrary finite numbers of risk classes, provided the different risk classes can be completely ordered (in terms of expected cost at any policy) and that their utility functions have the "single-crossing" property: For any two risk classes, and any two indifference curves, one for each class, there is at most one point where the curves cross. However, as the number of types increases, the existence of a Nash equilibrium becomes in some sense less likely. Indeed, Riley showed that in the limiting case of a continuum of types, there is never a Nash equilibrium of the sort described by RS.

Within the context of the one-policy-per-firm model, there have been two "resolutions" of the RS nonexistence problem. First, Wilson described a different solution concept, termed by him "anticipatory equilibrium," in which the usual Nash assumption that the strategies of

the other firms will not be affected by one's own action is replaced by the assumption that policies which become unprofitable in the face of entry will be dropped. In the single-policy-per-firm context, this is equivalent to the assumption that firms expect that rival firms on whom they inflict losses will go out of business. An equilibrium relative to these conjectures is a situation in which each firm makes zero profits (actuarial fairness due to free entry), and in which any profitable entry would become unprofitable once the reactions anticipated by the entrant have occurred. With this definition, equilibrium always exists. It coincides with the RS equilibrium when the latter exists, and otherwise represents a "pooling" equilibrium, in which several different types purchase the same policy.

It is important to realize that the Wilson equilibrium can also be described as a Nash equilibrium, albeit of a slightly different game. Suppose that we construct an extensive-form game corresponding to the story just given: Starting from some initial point, firms will offer policies, observe the results, and enter and exit, until a steady state is reached, at which point (in this timeless world) they collect their payoffs. In the RS story, it may well be the case that no branch of this tree (which we shall assume has perfect information) ever terminates; this corresponds to the absence of equilibrium.[1] If there was an RS equilibrium, then some branches would terminate, and they would do so in an "endgame" of the following stylized form:

... incumbents pick zero-profit policies	... entrants ... decide

It should be remarked that we can treat the possibility of infinite branches by assigning to them zero payoffs for all players.

In the Wilson model, the situation is the same, except that there is one more move of the incumbents, and it is of a special nature: They decide only whether or not to remain in the market. They cannot adjust

[1] We could equally well describe a simultaneous-move "endgame" between incumbent and entrant that did not have a pure-strategy equilibrium.

their policies, for if they did so we would essentially have returned to the RS world with the roles of incumbent and entrant reversed, and there would be no equilibrium. The stylized endgame is now:

... incumbents	... entrants	... incumbents
pick	pick	exit or stay

A way to describe this game in normal form is to specify a two-part strategy for the incumbents. A strategy now consists of two things: a policy to offer, and a reaction to any pair consisting of incumbent policies and entry policies. This reaction may assume two values: the previous policy (stay) or no policy (exit). Backwards induction will allow us to find the (perfect) equilibria, and it is evident that the reaction part of the incumbent strategy must correspond to the Wilson anticipations. It follows that the Wilson equilibrium is merely perfect equilibrium in a richer strategy space.

Another approach to the RS nonexistence problem is contained in a recent paper by Dasgupta and Maskin (DM), and represents the usual solution to the failure of Nash equilibria in pure strategies to exist: mixed strategy Nash equilibria. DM show that such equilibria always exist for the "adverse selection game," and also that these mixed-strategy equilibria possess some cross-subsidization properties. For our purposes, we merely note that they also represent Nash equilibria of a game with a wider strategy space than that allowed in the RS model.

All of these models retain the single-option requirement implicit in the constraint that each equilibrium policy offers zero expected profits. RS considered the possibility of offering multiple options, but found that it did not obviate the nonexistence problem. Indeed, as we shall see, it exacerbates the problem somewhat. Wilson did not consider multiple options, and it is natural to ask what the market would look like if firms offered multiple-option insurance, but retained anticipatory conjectures a la Wilson. This is the equilibrium concept that is the subject of this report. It too involves broadening the RS strategy space, but in a different way. Firms are now allowed to offer

coordinated portfolios of policies that may involve explicit cross-subsidization. This idea was first developed in another context by Miyazaki, and applied to insurance markets by Spence, although his definition differs slightly from ours in the case of many risk classes. This would seem to be a natural extension for two reasons. First, it is undeniably the case that cross-subsidization does exist (e.g., between Blue Cross High and Low Options) between distinct policies, and is not merely implicit between different risk classes purchasing the same policy as in the pooling equilibrium. Second, the fact that insurance offered to different risk classes introduces an externality between the profits of various policies suggests an analogy with the theory of joint products, according to which subsidization is entirely reasonable.

Spence's analysis was mainly devoted to showing that such policies result in a solution to a particular optimization problem, first identified in this context by RS as a benchmark of efficiency, and to examining other measures of efficiency. In order to obtain the existence of an equilibrium supporting this optimum, Spence was forced to resort to a modified version of the conjecture entertained by firms in Wilson's model: that firms will react to entry by dropping unprofitable policies. The simple statement that firms drop unprofitable policies is obviously inconsistent with subsidization, so some modification was necessary. In essence, players should anticipate that firms which experience losses as a result of entry will go out of business. This is precisely the conjecture we shall use. It is much simpler than an assumption based on dropping unprofitable policies could be if it were to allow subsidization, and it has certain conceptual advantages as well. If firms are to be allowed to drop only the unprofitable portions of their portfolios after entry, two questions naturally arise:

- Why did they not have this freedom before entry occurred?
- If they are allowed to alter (rather than delete) portions of their portfolios after entry, why should they not be allowed to alter the terms of the policies?

If we allow firms either of these freedoms, we return to the RS world, with one important difference: The domain of existence of RS equilibrium is even smaller than it was before, since the RS policies, if they do not solve the optimal-subsidy problem that we shall describe below, can always be attacked by a Pareto-superior, strictly profitable pair of mutually subsidizing policies.

As far as efficiency is concerned, the optimization problem that is solved by the Spence-Miyazaki (SM) equilibrium in the two-type case is that of maximizing the utility of the low-risk types subject to the constraints of free choice by each type, zero overall profits (free entry), and the additional constraint that the high-risk types receive at least their full fair insurance utility. The latter constraint rules out perverse subsidies from high- to low-risk types. This problem was first defined by RS in their discussion of the suboptimality of the RS separating equilibrium. We can compare this with the problem of maximizing the utility of the low-risk types subject to the first two constraints alone, which would be appropriate to a regulated insurance market with controlled entry. If we refer to the latter as full optimality and the former as constrained optimality, we may summarize the results for the two-type case as follows: When RS equilibrium exists in the single-policy world, it may or may not be constrained optimal; the Wilson equilibrium in the single-policy world is only constrained optimal if it is the RS equilibrium as well; and there is a unique population distribution for which either is fully optimal. With multiple-policy strategies, using Nash equilibrium in Game I, the RS equilibrium, when it exists, is always constrained optimal; the critical population distribution at which it ceases to exist is precisely the distribution for which it is fully optimal; and the policies that characterize it are the same as those characterizing constrained optimal RS equilibria. If we play according to Game II, we obtain an equilibrium that coincides with the Game I equilibrium when the latter exists; exists for all distributions; and is always fully optimal for those distributions having no Game I equilibrium. In the intermediate range between the distribution at which the RS equilibrium is fully optimal and that at which it ceases to exist in the single-policy world, the Game II equilibrium Pareto dominates the RS equilibrium.

OUTLINE OF THE STUDY

Section II below describes the optimal subsidy problem for the two-type case, and characterizes full and constrained optima. Section III addresses the strategic model in the two-type, multiple-option world, providing existence, uniqueness, and optimality results. Section IV discusses the perfectness properties of the equilibria defined in Sec. III; and Sec. V extends the analysis to the case of an arbitrary finite number of types. Finally, Sec. VI contains a discussion of the limitations of this model and suggestions for further extension.

II. OPTIMALITY IN A TWO-TYPE MODEL

We are investigating the following situation: There are n_1 agents of type 1 and n_2 agents of type 2. For each agent there are two possible states of nature, so the commodity space we deal with consists of pairs of state-contingent incomes, (a,b), where a is interpreted as the agent's income if an accident occurs and b his income if no accident occurs. All agents have the same initial endowment, which we denote (a_0,b_0), and the same strictly increasing, strictly concave, indirect utility function, U, for income. An agent of type i has a probability t_i of having an accident. For concreteness, we assume $t_1 > t_2$, so an agent of type 1 will be called a high-risk agent, while an agent of type 2 is a low-risk agent.

State-contingent income allocations (a,b) are evaluated according to the agents' von Neumann-Morgenstern expected utility functions:

$$(1) \qquad U_i(a,b) = t_i\, U(a) + (1 - t_i)U(b)$$

As a matter of convenience, we make the following definitions:

$$n = n_1/(n_1 + n_2)$$

$$t = nt_1 + (1 - n)t_2$$

$$x = ta_0 + (1 - t)b_0$$

Considerably greater generality is possible, in terms of differing endowments and differing utility functions, but it will not materially affect the results unless it ceases to be the case that representative utility curves of different types intersect in at most one place, and that at a point of intersection, the indifference curve corresponding to the higher-risk agent is flatter than that of the lower-risk agent.

We shall define three concepts of efficiency to use as benchmarks in this model. The first, which we denote E1, represents efficiency in a full-information world, where the only constraints are individual rationality and the requirement that the overall portfolio break even.

E1: Maximize $wU_1(a_1,b_1) + (1 - w)U_2(a_2,b_2)$ subject to:

(i) $U_i(a_i,b_i) \geq U_i(a_0,b_0)$, all i, and

(ii) $n(t_1a_1 + (1 - t_1)b_1) + (1 - n)(t_2a_2 + (1 - t_2)b_2) \leq x$

E2 represents the sort of efficiency one might hope for in a world of asymmetric information; it differs from E1 by the presence of an additional constraint which specifies that no type prefers, over its own policy, the policy designed for the other type. In practice, this constraint will only be binding on the high-risk types.

E2: Maximize $wU_1(a_1,b_1) + (1 - w)U_2(a_2,b_2)$ subject to:

(i) $U_i(a_i,b_i) \geq U_i(a_0,b_0)$, all i,

(ii) $n(t_1a_1 + (1 - t_1)b_1) + (1 - n)(t_2a_2 + (1 - t_2)b_2) \leq x$, and

(iii) $U_1(a_1,b_1) \geq U_1(a_2,b_2)$

Finally, E3 is the efficiency concept appropriate to a competitive world of asymmetric information. In such a world it is easy to show (as we do formally in Sec. III) that the high-risk agents must achieve at least their full fair-insurance utility $U(t_1a_0 + (1 - t_1)b_0)$. If not, then policies exist that would be strictly preferred by the high-risk agents and would earn positive profits regardless of which group or groups purchased them. This is the optimization problem defined by RS (II.3) and used in their discussion of efficiency.

E3: Maximize $wU_1(a_1,b_1) + (1 - w)U_2(a_2,b_2)$ subject to:

(i) $U_i(a_i,b_i) \geq U_i(a_0,b_0)$, all i,

(ii) $n(t_1a_1 + (1 - t_1)b_1) + (1 - n)(t_2a_2 + (1 - t_2)b_2) \leq x$,

(iii) $U_1(a_1,b_1) \geq U_1(a_2,b_2)$, and

(iv) $U_1(a_1,b_1) \geq U(t_1a_0 + (1 - t_1)b_0)$

Definition: The portfolio $[(a_1,b_1),(a_2,b_2)]$ is E_i-optimal, for i = 1,2,3, if it solves problem E_i.

It is evident that these efficiency concepts are increasingly restrictive. It is also true that efficiency implies that the high-risk agents receive full insurance.

Proposition 2.1: Let $[(a_1,b_1),(a_2,b_2)]$ be E_i-optimal.

Then $a_1 = b_1$ and, if i = 1, $a_2 = b_2$.

PROOF: The Lagrangeans corresponding to each problem have the following symmetry property: Whenever a term involving a_1 appears, it is in the form $t_1f(a_1)$ for some function f. There is always a corresponding term of the form $(1 - t_1)f(b_1)$ and vice versa. This means that differentiating the chosen Lagrangean w.r.t. a_1 gives the same first-order condition as differentiation w.r.t. b_1; and since all terms involving $a_1(b_1)$ are either linear in a_1 or are multiples of $U(a_1)(U(b_1))$, we conclude that $U'(a_1) = U'(b_1)$ and thus by concavity of U that $a_1 = b_1$. The same argument shows that $a_2 = b_2$ in problem E1. QED

In light of this proposition, we shall maintain the following change of notation for the balance of this section:

$$a_1 = b_1 = c$$
$$a_2 = a$$
$$b_2 = b$$

Examination also shows that unless $t_1 = t_2$, E2 and E3 imply $a_2 \neq b_2$, so that E2 and E3 optimality give smaller values of the objective function than E1 optimality.

It may be useful to give a geometric construction showing the solutions to the optimization problems. There are many representations of insurance contracts; we find it most useful to work directly in terms of the state-contingent net allocations. Thus, the commodity space we shall work with is a subset of R^2. In fact, with initial allocation (a_0, b_0), it is unnecessary to consider any allocation with a coordinate greater than $a_0 + b_0$. Figure 1 represents the initial situation, and the solution to problem E1.

The diagram shows three "breakeven lines": The set of policies that would break even if purchased by the high-risk types is denoted be1, and defined by $t_1 a + (1 - t_1)b = t_1 a_0 + (1 - t_1)b_0$; the set of policies that would break even if purchased by the low-risk types is denoted be2, and defined by $t_2 a + (1 - t_2)b = t_2 a_0 + (1 - t_2)b_0$; the set of policies that would break even if purchased by (a random sample of) the "market" population is denoted beM and defined by $ta + (1 - t)b = ta_0 + (1 - t)b_0$. The E1-optimal policy for type i is located at the

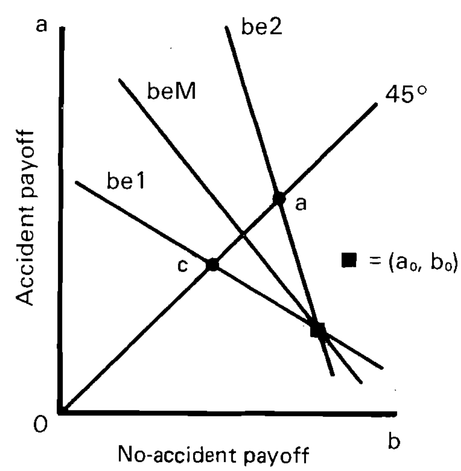

Fig. 1 -- The E1-optimal policies

intersection of bei with the full-insurance line defined by a = b; this is where U_i is maximized on bei.

The E1 optima are achieved at unique policies. Since they involve full insurance for both types, the E1 problem can be written:

$$\text{maximize } wU(c) + (1 - w)U(a) \text{ s.t.}$$
$$nc + (1 - n)a = x,$$

where a is the full-insurance outcome for the low-risk players. The uniqueness of a and c follows by concavity.

When we move to the asymmetric-information setting, the separation constraint (iii) comes into play. First, we can define the E2-optimal contracts for the two types that lie on their respective breakeven lines. As RS showed, this pair of contracts, which comprise their separating contracts, consist of the E1-optimal high-risk policy, denoted by C, together with the policy on be2 that maximizes U_2 subject to separation. This latter is denoted R in Fig. 2.

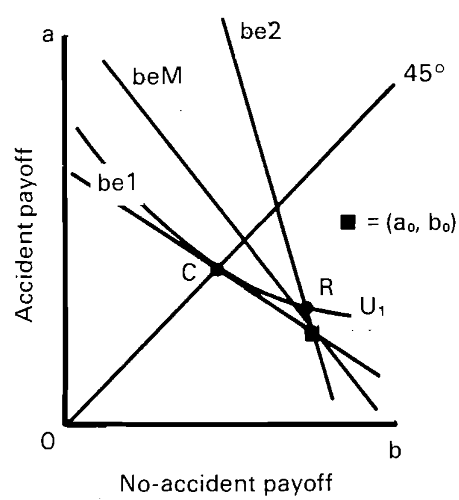

Fig. 2 -- The "separating contracts"

This pair of contracts can be used as an initial feasible solution to problem E2. We shall now construct the locus of all feasible solutions to E2 that might be optimal. This entails the constraints (i)-(iii) above, together with $a_1 = b_1$, in light of Proposition 2.1. The locus we want is that of contracts for the low-risk types which would make zero overall profits when offered in conjunction with a riskless contract that is equally ranked by the high-risk types. Clearly, R is on this locus, when offered in conjunction with C. Another point on this locus is the intersection of the market-odds line beM with the full-insurance locus.

Now pick any point B on beM. If this were purchased by both types, it would make zero profits. If the high-risk types were sold any contract on a line through B parallel to be1, and the low-risk types were sold any contract on a line through B parallel to be2, that combination would also make zero profits. (See Fig. 3.)

Conversely, given any zero-profit pair of contracts $[(a_1,b_1),(a_2,b_2)]$, if we draw a line parallel to be1 through (a_1,b_1), and a line parallel to be2 through (a_2,b_2), the lines will intersect beM at the same point.

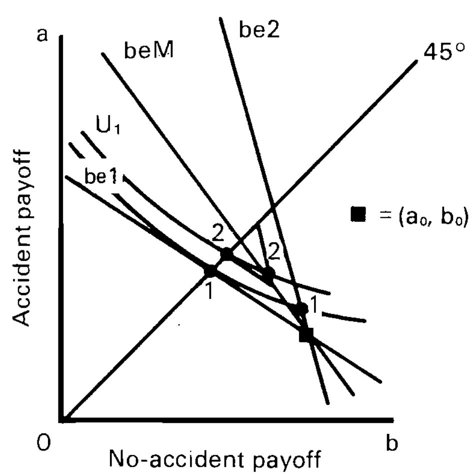

Fig. 3 -- Some zero-profit pairs

In light of the individual rationality constraint (i) above, we shall start the construction using that riskless policy that gives the high-risk types their initial utility $U_1(a_0,b_0)$. This is denoted by C_0 in Fig. 4. In conjunction with this policy, the low-risk types must be offered a policy along the line labeled L_0. By quasi-concavity of U, there are at most two points on L_0 that satisfy the separation constraint (iii). Of these, the "lower" one, with $a < b$, offers a higher utility to the low-risk types, in view of their steeper indifference curves and the single-crossing property. This leads to the selection of the policy labeled D for them.

The other end of the feasible locus is the point A. It lies on the locus, and if the high-risk types are offered any riskless policy preferred to A, the policy offered to the low-risk types lies on a line parallel to be2 that intersects beM where $a > b$. This, combined with the steeper slope of the low-risk indifference curves, means that they would prefer the policy offered to the high-risk types. Since this policy lies above beM, its purchase by both types is incompatible with zero profits.

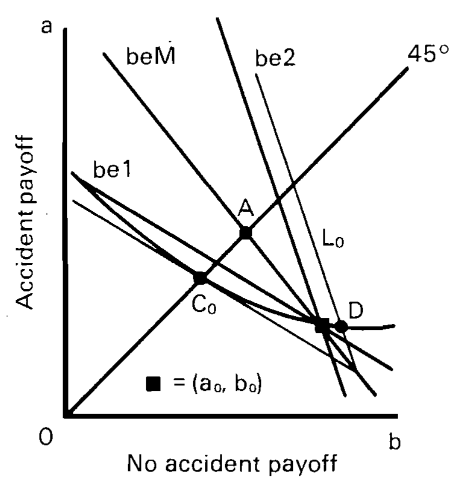

Fig. 4 -- The ends of the zero-profit locus

We can therefore generate the entire zero-profit locus in the following manner: Beginning with any point E on the full-insurance line between C_0 and A, we first find the point F where the line through E parallel to be1 intersects beM. Let us denote by L(E) the line through F parallel to be2; the policy for the low-risk types, G, is located at the (lower) intersection of L(E) and the indifference curve of the high-risk types passing through E.

The entire locus, labelled Π_o, is shown in Fig. 5.

Several things should be clear from this construction. First, if we limit our attention to points below A as argued above, then the separation constraint only affects the high-risk types. Second, the locus lies above beM and intersects be2 only at R. Third, the slope of the locus can be determined directly from the constraints.

Using Proposition 2.1, we get:

$$\text{(ii)}\ nc + (1 - n)(t_2 a + (1 - t_2)b) = x,$$
$$\text{(iii)}\ U(c) = t_1 U(a) + (1 - t_1)U(b)$$

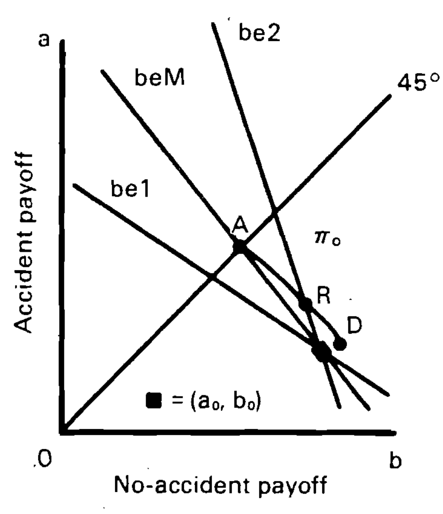

Fig. 5 -- The zero-profit locus

Total differentiation gives:

$$\text{(ii')}\quad ndc + (1 - n)(t_2 da + (1 - t_2)db) = 0$$
$$\text{(iii')}\quad U'(c)dc = t_1 U'(a)da + (1 - t_1)U'(b)db$$

The slope of the locus is therefore:

$$\text{(2)}\qquad da/db = - \frac{[n(1-t_1)U'(b) + (1-n)(1-t_2)U'(c)]}{[n(t_1)U'(a) + (1-n)(t_2)U'(c)]}$$

Since the market will direct our attention to the special case w = 0, we shall concentrate on it for the balance of the section. There are several possibilities for the solution of the optimization problems E2 and E3. Under certain circumstances, the solution will be an interior maximum of the utility of the low-risk types among the contracts on the zero-profit set. If such an interior maximum occurs between A and R, then it is a solution to both the E2 and E3 problems (see Fig. 6).

If such an interior maximum falls between R and D, then it is E2-optimal but not E3-optimal. The reason is that it violates constraint (iv). More interesting is the observation that such an optimum involves "perverse subsidies" in the sense that the insurance firm would earn profits on its high-risk clientele which would pay for the losses it suffered on the low-risk policies. This is the reason why

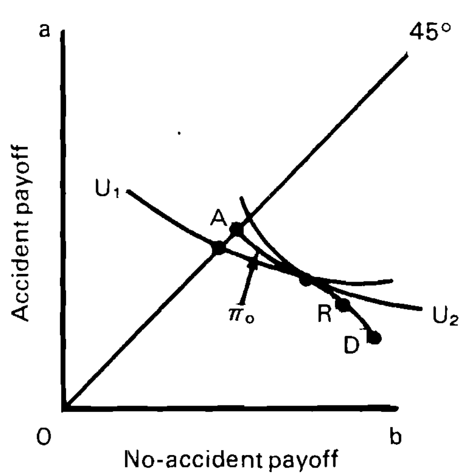

Fig. 6 -- An interior E2 and E3 optimum

such an optimum would not be expected to prevail in a market. In such a case the E3 optimum would be R. This is shown in Fig. 7.

A corner solution at A is impossible, since the zero-profit locus and the market-odds line beM are tangent at A. However, since this common slope is flatter than that of be2, it is flatter than a low-risk indifference curve, which has slope $-(1 - t_1)/t_1$ along the full-insurance line.

It is possible to continue the zero-profit locus beyond D, and in principle one might expect the possibility of optima located beyond D. However, this would violate the individual-rationality constraint (i) for the high-risk types.

In case there is an interior maximum, a useful necessary condition can be obtained by equating the slope of a low-risk indifference curve to that of the zero-profit locus found above. This gives:

$$\frac{(1-n)t_2(1-t_2)}{n(t_1-t_2)} = U'(c)\left[\frac{1}{U'(a)} - \frac{1}{U'(b)}\right] \tag{3}$$

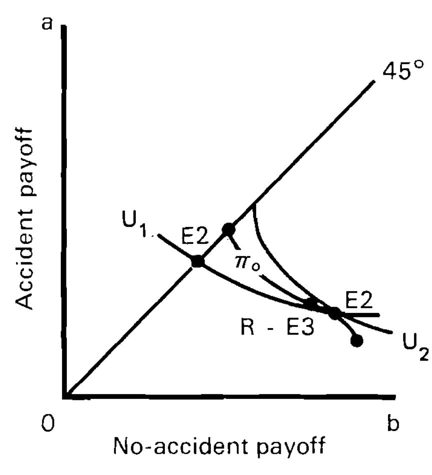

Fig. 7 -- Separate E2 and E3 optima

As the population parameter n varies, the positions of the optima will vary as well. We shall demonstrate that the solution to E2 is unique and varies continuously with n, and that there is a well-defined "critical value" n^* such that the solution to E2 solves E3 if $n < n^*$. On the other hand, if $n > n^*$, the solution to E3 will be R, while the solution to E2 will involve perverse subsidies, and will lie between R and D.

Proposition 2.2: The solution to E2 is unique in the two-types case, if U is twice continuously differentiable and strictly concave. It is also a continuous function of the population parameter n.

PROOF: We can write E2 as:

$$\max t_2 U(a) + (1 - t_2) U(b) \text{ s.t.}$$
$$U[(x - (1 - n)(t_2 a + (1 - t_2)b))/n] \geq t_1 U(a) + (1 - t_1) U(b)$$

The IOC's defining the solutions to this give the optimality condition, Eq. (3). If multiple solutions to Eq. (3) exist, then they must give the same utility to the type 2 agents (we are only interested in the uniqueness of the global solution to the problem, not whether other local solutions might exist).

Thus, there must exist (a,b) and (a',b') satisfying Eq. (3) and s.t.

$$(4) \qquad t_2 U(a) + (1 - t_2) U(b) = t_2 U(a') + (1 - t_2) U(b')$$

In particular, if we use this condition to write b implicitly as a function of a, then we want to show that the LHS of the above equation is monotone in a. Let us define the following functions:

$$(5) \qquad \begin{aligned} V(x) &= 1/U'(x) \\ W(x) &= U''(x)/(U'(x))^2 \end{aligned}$$

We can now write the LHS of Eq. (3) as $U'(c)[V(a) - V(b)]$, and its derivative (w.r.t. a) as:

$$(6) \qquad U'(c)[V'(a) - V'(b)b'] + U''(c)c'[V(a) - V(b)]$$

However,

$$(7) \qquad V'(x) = -W(x)$$

so the quantity in (6) can be written:

$$(8) \qquad U'(c)[W(b)b' - W(a)] + U''(c)c'[V(a) - V(b)]$$

We can evaluate b' and c' by taking the total derivatives of Eq. (4) and of the constraint in E2:

$$(9) \qquad b' = -[t_2/(1 - t_2)][V(b)/V(a)]$$

$$(10) \qquad c' = [(t_1 - t_2)/(1 - t_2)][V(c)/V(a)]$$

Substitution of Eqs. (9) and (10) in Expression (8) and division by $U'(c) > 0$ gives:

$$(11) \qquad W(c)\left[\frac{[(t_1-t_2)]}{[(1-t_2)]}\right][1 - V(b)/V(a)] - W(b)\left[\frac{t_2}{(1-t_2)}\right][V(b)/V(a)] - W(a)$$

But strict concavity implies that W(x) is uniformly negative, and moreover since $a < b$, we have $V(b) > V(a)$, so Expression (11) is positive in each term. Hence the RHS of Eq. (3) is strictly monotone in a and there exists but a single global solution to E2. Moreover, the LHS of Eq. (3) is monotone-decreasing and continuous in n, so this unique solution is also continuous in n. QED

In solving E3, if constraint (iv) is binding, we are at R, and we may specify the contracts completely using the following versions of the constraints:

(12) $c(R) = t_1 a_0 + (1 - t_1) b_0$

(13) $t_2(a(R) - a_0) = (1 - t_2)(b_0 - b(R))$

(14) $U(c(R)) = t_1 U(a(R)) + (1 - t_1) U(b(R))$

None of these conditions depends on n, so (a(R),b(R),c(R)) is independent of n, as is the RHS of Eq. (3) evaluated at the R-contracts. The LHS is monotone-decreasing and continuous in n, so that n* defined by:

$$(15) \qquad \frac{(1-n^*)t_2(1-t_2)}{n^*(t_1-t_2)} = U'(c(R)) \left[\frac{1}{U'(a(R))} - \frac{1}{U'(b(R))}\right]$$

is unique and has the properties claimed above.

If U is not strictly concave, there may be multiple optima, and they may vary discontinuously with n. In Secs. III and IV, we shall assume that the E2-optimal policy is unique.

Proposition 5.2 below provides a generalization of Proposition 2.2 above, with a simpler proof.

III. EQUILIBRIUM IN A TWO-TYPE WORLD

As noted in the Introduction, the nature of equilibrium in markets with asymmetric information is heavily dependent on the conjectures entertained by firms about the reactions of their rivals and on the strategies employed by firms. All of the models under discussion involve a free-entry condition, which means that whatever the equilibrium set of policies may be, it will be characterized by expected profits of zero for each firm. Under the one-policy-per-firm models of RS and W, this in turn implies that each policy must break even for the group it is intended to serve. In multiple-option models, this means that the portfolio of policies offered by a firm, while it may involve some cross-subsidization, breaks even overall.

In essence, these models are models of equilibrium with externalities; since the consumers are not assumed to behave strategically, the asymmetric information merely serves as an externality linking various policies on the market, and multiple-option policies provide a means of internalizing the externality.

The strategic assumptions may be phrased in various ways, but they can be reduced (for the two-type case) to two: no reaction by rivals in terms of the policies offered (which we call Nash conjectures); and no reaction by rivals in terms of strategies, where a strategy may make policy offerings contingent on nonnegative profits (which we call anticipatory conjectures after W). Either type of conjecture can be used to define a conjectural equilibrium in the sense that no firm conjectures that it can increase its profits by a unilateral change of strategy. It is also clear that for every equilibrium in strategies that say, "offer policies p," there is an equivalent equilibrium in strategies that say "offer policies p as long as they earn nonnegative expected profits; otherwise offer no policies."

In the present model, we suppose that individual firms are free to offer as many policies as they wish. A strategy for a firm is thus a set $\{(a_i,b_i)\}$ of policies (state-contingent allocations), together with a statement as to which if any of these policies will be offered should

they prove to be unprofitable. The nature of this "second" part of the strategy threatens to become unmanageable; in general, we are asking for something like a strategy which describes the set of policies that will be offered as a function of the spectrum of policies offered by the other firms.

While this may be a fruitful way to think of the strategy space, it poses analytical difficulties. Let the space of policies be denoted P; it is a convex compact subset of R^2. If P~ is the power set of P, then a strategy can be viewed as a function:

$$s(i):P\sim \longrightarrow P\sim$$

To each n-tuple of strategies, s, we would like to assign an outcome R(s) belonging to P~xP~x ... P~. Denoting this as R(s) = R(1,s), ..., R(n,s), and writing R(-i,s) as the union of all R(j,s) except R(i,s), we require that:

$$R(i,s) = s(i)(R(-i,s))$$

for each i.

To each outcome R = R(1), ..., R(n) we can assign payoffs, so that in principle it is easy to write the normal form of this game. We shall denote the payoff function thus obtained as H(s). In what follows, we shall confine our attention to some particular special cases of this setup.

Game I: The Rothschild-Stiglitz game

In this version, each s(i) is a single element of P; moreover, it does not depend on the policies offered by the other firms. Let H(i,(ai,bi)), $(a_{-i,}b_{-i})$ be the profits of firm i when it offers (a_i,b_i) and its rivals offer (a_{-i},b_{-i}).

Game II: The Wilson game

In this version, each s(i) is again a single element of P. It depends on the policies chosen by other firms in the following manner: Each s(i) has the form:

$$\begin{aligned} s(i)(R(-i,s)) &= (a_i,b_i) \text{ iff } H(i,(a_i,b_i),R(-i,s)) \geq 0 \;; \\ &= (a_0,b_0) \text{ otherwise.} \end{aligned}$$

Game III: The multiple-option Nash game

In this version, each s(i) is a subset of P; however, it does not depend on the policies chosen by the other firms.

Game IV: The multiple-option anticipatory game

In this version, each s(i) is a subset of P; the dependence on the policies chosen by the other firms is as follows:

$$\begin{aligned} s(i)(R(-i,s)) &= \{(a_i,b_i)\} \text{ iff } H(i,\{(a_i,b_i)\},R(-i,s)) \geq 0; \\ &= \{(a_0,b_0)\} \text{ otherwise.} \end{aligned}$$

In other words, each firm chooses a fixed set of policies that it will offer if they make nonnegative profits, and offers the null policy otherwise.

These games differ only in the strategy spaces of the players; the definition of equilibrium is the same for each game:

Definition 3.1: An equilibrium is an n-tuple $s^* = s^*(1), \ldots, s^*(n)$ of strategies with the property that for each i, and each s(i),

$$H(i,s^*) \geq H(i,s(i),s^*(-i))$$

where $s(-i) = s(1), \ldots, s(i-1), s(i+1), \ldots, s(n)$.

Although all these games have compact convex spaces of pure strategies, the payoff function is not continuous, so there is no guarantee that an equilibrium in pure strategies exists. In fact, we shall find that a pure-strategy equilibrium always exists for Games II and IV, but does not always exist for Games I and III.

There are some simple relations between the various games. An equilibrium in either of the games with Nash conjectures (I or III) gives rise to an equilibrium with an identical outcome in the corresponding game with anticipatory conjectures (II or IV, respectively). Moreover, since the strategy spaces for the single-option games (I and II) are subsets of the multiple-option strategy spaces for the games with corresponding conjectures (III and IV, respectively), an equilibrium in the latter games that happens to involve single-option policies will be an equilibrium for the former games. However, equilibria of the single-option games are not necessarily equilibria of the corresponding multiple-option games. We shall find that there is a close connection between this question and the optimality properties of the single-option equilibria.

Throughout the balance of this section, we shall focus on the influence of the population parameter, n, on the existence and optimality properties of equilibria. We begin by defining the conditions under which equilibria of Games I and III will exist, and by showing that equilibria for Games II and IV always exist.

Proposition 3.2: There exist $n^{\sim} < n^{*}$ with the following properties:

(i) For all $n \geq n^{\sim}$, there exists a pure-strategy equilibrium of Game I; if $n < n^{\sim}$, there is no pure-strategy equilibrium.

(ii) For all $n \geq n^{*}$, there exists a pure-strategy equilibrium of Game III; if $n < n^{*}$, there is no pure-strategy equilibrium.

(iii) For $n \geq n^{*}$, the sets of equilibrium policies for Games I and III are the same.

PROOF: Property (i) is due to RS. The equilibrium set of policies, when it exists, consists of two distinct policies: The high-risk agents are offered full fair insurance, and the low-risk agents receive the best fair insurance that would not attract the high-risk types. To be precise, the high-risk types get:

$$a_1 = b_1 = c(R) = t_1 a_0 + (1 - t_1) b_0$$

and the low-risk types' policy $(a(R),b(R))$ satisfies

(1) $$U(c(R)) = t_1U(a(R)) + (1 - t_1)U(b(R))$$

(2) $$t_2(a(R) - a_0) + (1 - t_2)(b(R) - b_0) = 0$$

(3) $$a(R) \leq b(R)$$

where the first condition is the separation condition; the second condition is the fair-odds condition; and the third condition, together with the fact that

(4) $$t_1 > t_2,$$

guarantees that the chosen policy offers greater utility to the type 2 agents than any other solution to (1) and (2).

The failure of equilibrium to exist if $n > n\sim$ stems from the fact that a "pooling" policy located on (or below) beM may offer higher utility to the type 2 agents than $(a(R),b(R))$ does. In this case, such a policy will attract the type 1 agents as well. The existence condition is therefore that beM lie below the type 2 indifference curve through $(a(R),b(R))$. This means that:

(5) $$\begin{aligned} V(n) &= \max\, t_2U(a) + (1 - t_2)U[t(a_0 - a)/(1 - t) + b_0] \\ &\leq t_2U(a(R)) + (1 - t_2)U(b(R)) \end{aligned}$$

Since $t = n_{t1} + (1 - n)t_2$ is monotone-increasing in n, $V(n)$ is strictly monotone-decreasing in n (as the market line beM approaches the type 1 fair-odds line be1, the best type 2 utility along that line decreases).

Therefore, there is a single, well-defined $n\sim$ at which Eq. (5) holds with equality. For $n < n\sim$, Eq. (5) does not hold, and the policy $(a(P,n),b(P,n))$ that solves the problem on the LHS of Eq. (5) will be strictly preferred to the RS policy by both types, and will break even.

However, this policy cannot be an equilibrium policy (unless $n = 0$), by virtue of the fact that the low-risk agents have steeper indifference curves than the high-risk agents). Figure 8 shows that there will exist a policy preferred to (a(P,n),b(P,n)) by only the low-risk agents. Since the market fair-odds line is less steep than bel, this policy can be chosen below bel, and will therefore be profitable.

Property (ii): The n^* referred to is that defined in Sec. II; it is the point at which the separating contracts (c(R),c(R)),(a(R),b(R)) are E2-optimal. First we shall show that $n^* > n\sim$. Then we shall show that $n > n^*$ implies that (c(R),c(R)),(a(R),b(R)) are equilibrium contracts. Finally, we shall show that for $n\sim < n < n^*$, there are no pure-strategy equilibrium contracts.

To see that $n^* > n\sim$ it is useful to recast their definitions. Let us assume that the E2-optimal policies are unique (for example, let U be smooth and strictly concave), and denote them (as functions of n) by

$(c^*(n),c^*(n))$ for the high-risk players, and

$(a^*(n),b^*(n))$ for the low-risk players.

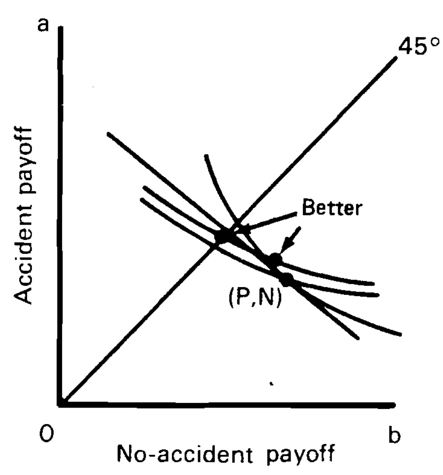

Fig. 8 -- A contract that beats (a(P,n),b(P,n))

Recall that n^* is defined by:

(6) $[(c^*(n^*), c^*(n^*)), (a^*(n^*), b^*(n^*))] = [(c(R), c(R)), (a(R), b(R))]$

and that $n\sim$ is defined by:

(7) $V(n\sim) = t_2 U(a(R)) + (1 - t_2)U(b(R))$

By definition, $n^* > n\sim$ if there exist a subsidizing pair of policies that is Pareto-preferred to both $(a(P,n\sim), b(P,n\sim))$ and $(c(R),c(R)),(a(R),b(R))$ and that earns strictly positive profits at $n\sim$. By definition,

(8) $t_2 U(a(P,n\sim)) + (1 - t_2)U(b(P,n\sim)) = t_2 U(a(R)) + (1 - t_2)U(b(R))$

and since $t_1 > t_2$,

(9) $t_1 U(a(P,n\sim)) + (1 - t_1)U(b(P,n\sim)) > U(c(R))$

Therefore, consider the pair of policies $[(c',c'), a(P,n\sim), b(P,n\sim)]$ defined by:

(10) $U(c') = t_1 U(a(P,n\sim)) + (1 - t_1)U(b(P,n\sim))$

The losses when the type 1 agents all purchase the policy (c',c') are strictly the minimum losses over all policies giving them a utility of at least $t_1 U(a(P,n\sim)) + (1 - t_1)U(b(P,n\sim))$. Therefore, by continuity U and of linear functions, there exist positive numbers ε and δ such that:

(11) $n[c' + \delta] + (1 - n)[t_1 a(P,n\sim) + (1 - t_1)b(P,n\sim) + \varepsilon] < x$

(12) $U(c' + \delta) > t_1 U(a(P,n\sim)) + (1 - t_1)U(b(P,n\sim))$, and

(13) $t_2 U(a(P,n\sim) + \varepsilon) + (1 - t_2)U(b(P,n\sim) + \varepsilon) > V(n\sim)$

This shows that $n^* > n\sim$.

Now suppose that $n \geq n^*$. In this case, we know that $(c(R),c(R)),(a(R),b(R))$ is E3-optimal. Moreover, we know that it is an equilibrium of Game I, which means that any policy that is profitable when $(c(R),c(R)),(a(R),b(R))$ is on the market must be a separating pair of policies. However, any such pair must necessarily involve perverse subsidies, which in turn means that it must offer the high-risk agents lower utility than $U(c(R))$. The result of offering such a pair is that only the low-risk agents would purchase it, resulting in losses for the potential entrant.

In case $n^* > n \geq n\sim$ the separating policies $(c(R),c(R)),(a(R),b(R))$ are not E2 (which coincides with E3) optimal. Since there exists a Pareto-preferred pair of subsidizing policies with zero profits, a continuity argument such as that used above shows that neither the separating policies nor any other pair of policies that are not E2-optimal can be equilibrium policies. However, the E2-optimal policies are also not equilibrium policies, for the same reason that the pooling policies were not equilibrium policies in Game I. At the policy offered to the low-risk agents, their indifference curves are steeper than the indifference curve of the high-risk agents, which runs through both their policy $(c^*(n),c^*(n))$ and through the low-risk policy $(a^*(n),b^*(n))$ by virtue of the separation constraint. It follows that there exist policies to the "southeast" of $(a^*(n),b^*(n))$ that would attract only the low-risk agents. Such policies, one of which is shown in Fig. 9, must be profitable to offer unless $(a^*(n),b^*(n))$ lies on or above bel, which contradicts $n < n^*$. For $n < n\sim$ the nonexistence argument is the same as for Game I. QED

In the course of this proof, we have obtained some efficiency results that are worth noting as a Corollary.

Corollary 3.3: The pure-strategy equilibrium of Game I is E3-optimal iff $n \geq n^*$. It is E2-optimal only if $n = n^*$. The pure-strategy equilibrium of Game III is E3-optimal iff it exists, and is E2-optimal iff $n = n^*$.

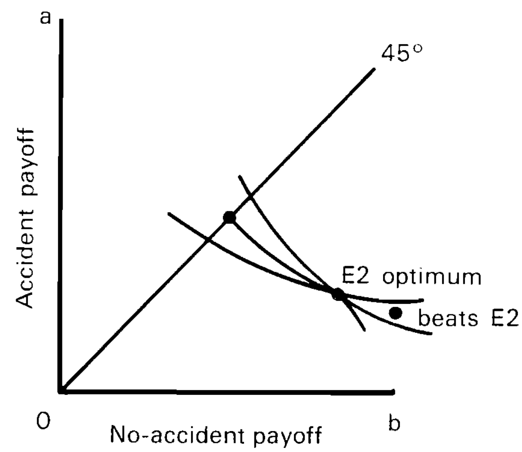

Fig. 9 -- A contract that beats $(a^*(n),b^*(n))$

When we turn to anticipatory equilibrium, the nonexistence problem disappears, but the efficiency properties become more interesting. Wilson has shown that for Game II, the pure-strategy equilibrium policies are given by $(c(R),c(R)),(a(R),b(R))$ when $n \leq n^\sim$ and are equal to $(a(P,n),b(P,n))$ for both players otherwise. For Game IV we have:

Proposition 3.4: Game IV always has a pure-strategy equilibrium. The policies offered are equal to $[(c(R),c(R)),(a(R),b(R))]$ for $n \leq n^*$, and are equal to $[(c^*(n),c^*(n)),(a^*(n),b^*(n))]$ otherwise.

PROOF: For $n \leq n^*$, we have seen that $[(c(R),c(R)),(a(R),b(R))]$ is an equilibrium of Game III. Moreover, since both policies break even, it is impossible for an entrant to inflict losses on an incumbent, so that entry via "loss leaders" is precluded. (In a two-type market, one can inflict losses on a firm offering a subsidizing pair of contracts only by attracting that firm's "winners"; if this is not profitable for the entrant, he must also be offering a contract for the incumbent's "losers" that will earn him positive profits once the incumbent is out of the market. However, such a reversal can only be profitable if $n \leq n^*$, in which case no losses can be inflicted.)

For $n > n^*$, we must show that $[(c^*(n),c^*(n)),(a^*(n),b^*(n))]$ is a set of policies that could be offered in equilibrium. First, since there does not exist a Pareto-improving pair of policies that offers nonnegative profits, it follows that any pair of policies capable of earning profits if it serves the entire market will not be purchased by at least one type when it is first offered. In order to induce the incumbent firm to exit, such an attacking pair of policies must attract the incumbent's low-risk customers. In order to do so, the low-risk customers must be offered a policy above the zero-profit locus derived in Sec. II. But then there is no policy for the high-risk types to purchase that they would prefer and that makes nonnegative overall profits in combination with the low-risk policy. QED

At this point, it is worth remarking that the above results do not depend on the uniqueness of the E2-optimal policy. In case there exist multiple solutions to E2, they must be Pareto-ranked, offering the same utility to the low-risk types and increasing utility to the high-risk types. In this case, the power of the competitive market to select optima is enhanced, in that it will select the Pareto-superior solution to E2. To see this, suppose that any of the other E2-optimal pairs of policies was being offered. Offering the low-risk agents a policy slightly better than that which they would purchase at the Pareto-superior solution to E2 and offering the high-risk types a policy slightly worse than the policy they would purchase at the Pareto-superior E2 optimum will result in attracting both types and earning positive profits. This is shown in Fig. 10.

The equilibrium argument of Proposition 3.4 works equally well when the attacking policy is a pooling policy. We can again sum up the efficiency results for the anticipatory games in a Corollary.

Corollary 3.5:
The pure-strategy equilibrium of Game II is:

E2-optimal if $n = n^*$;
E3-optimal if $n \leq n^*$; and
neither E2- nor E3-optimal if $n > n^*$.

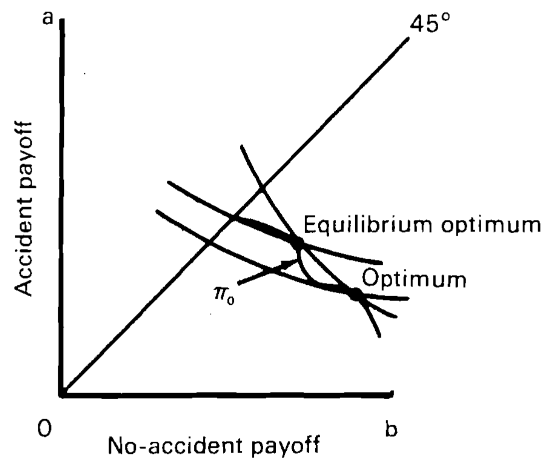

Fig. 10 -- Equilibrium selection among multiple E2 optima

The pure-strategy equilibrium of Game IV is:

E2-optimal if $n \geq n^*$; and
E3-optimal if $n \leq n^*$.

The results of this section indicate that the availability of multiple-option insurance will change the nature of market equilibrium. If firms entertain Nash conjectures, then the set of populations for which the RS equilibrium exists shrinks, since any suboptimal (in the E2 sense) RS separating policy can be attacked by a Pareto-superior, strictly profitable, pair of policies. This will in turn be vulnerable to attack by policies that siphon off the profitable low-risk agents.

On the other hand, if firms entertain anticipatory conjectures, the use of multiple-option policies serves to enhance the welfare properties of the set of policies offered in equilibrium. For the range of populations where the single-option anticipatory equilibrium separates types, the use of multiple options results in an unambiguous reduction in the externality due to asymmetric information, with the gains being

divided between the two types. For the range of populations where the single-option anticipatory equilibrium pools different types, the use of multiple options will increase the utility of the low-risk types but may decrease the utility of the high-risk types.

There is a sense in which the Nash and anticipatory conjectures form opposite ends of a continuum. Suppose that the game were to be repeated indefinitely, but that strategies could be selected only once. After a firm opened its doors, the only modification of its policies it could make would be to go out of business forever. If firms in such a game (which is an extreme example of what Marschak and Selten (1978) call an "inertia supergame") are completely myopic, in the sense that they completely discount the future, the Nash equilibria of this game will be equilibria of Games I and III. On the other hand, if firms evaluate their payoffs according to the limiting (undiscounted) average, then the equilibria of the repeated game will be equilibria of Games II or IV (depending on whether firms can offer multiple options). Seen in this way, it seems entirely reasonable that the true state of affairs lies somewhere between the extreme nonexistence of the myopic (Nash) conjectures and the extremely long-range view embedded in anticipatory conjectures. In any event, it seems that multiple options must form part of the model, since the existence of subsidization is indisputable.

IV. PERFECT EQUILIBRIUM

In addition to asking whether a particular set of policies would be offered as a result of market equilibrium, it is useful to have some information as to the sensitivity of such predictions to slight mistakes on the part of the players. This raises the issue of perfectness, since a perfect equilibrium is in some sense one that is immune to such small mistakes.

There are several different ways to define perfect equilibrium, and two of them have some application here. First, if a strategy is defined over more than one stage, or opportunity to play, then the latter prescriptions of that strategy are contingent on the earlier choices of all players. For example, in Game II or Game IV, the specification that a firm faced with losses on its aggregate portfolio would withdraw from the market gives an example of such a contingent prescription or threat. In such a game, one (relatively weak) perfectness requirement would have it that any such threat must be credible in the sense of being the most profitable available action should the circumstances leading to the threat actually occur.

This can be seen as requiring robustness to the small mistakes that lead to the invocation of the threat. If an opponent enters by mistake with a set of policies that inflict losses on the incumbent, it must be more profitable for the incumbent to exit than to carry on (we persist in the stipulation that no firm is allowed to modify an existing policy). It is fairly easy to see that the threat implicit in the strategies of Games II and IV is credible under this view. Indeed, the threat to continue to offer losing policies implicit in Games I and III is incredible. In the two-type case, it is also not a very intelligent threat, since there any entry that is eventually profitable is also immediately profitable.

A more subtle perfectness condition emerges when we study the normal form of the game, either by ignoring stages or by collapsing them into single "strategies." An equilibrium is perfect in the normal form iff it is the limit of a sequence of ε-perfect equilibria as ε

approaches 0. An ε-perfect equilibrium is a completely mixed-strategy n-tuple where no player uses a pure strategy that does not maximize his payoff with probability more than ε. Such strategies are called inferior strategies.

A strategy is dominated if there is another strategy that pays at least as much no matter what one's opponents choose, and pays strictly more for at least one choice by the opponents. If the opponents are playing completely mixed strategies, a dominated strategy is always inferior. Therefore, dominated strategies are never used in perfect equilibria.

If there are only two players in the game, this characterization is exact: An equilibrium is perfect iff it gives positive probability only to undominated strategies. If there are more than two players, it is possible for there to exist equilibria in undominated strategies that fail to be perfect. In this game, however, the strategies of the other players affect any one player's payoff only in the aggregate; what matters is only the set of policies offered by competing firms. Therefore, the sets of strategies attainable with independent and correlated play by the other players are the same, and the trembling-hand perfect equilibria are the undominated-strategy equilibria.

A vital question in seeing whether the equilibria of the previous section are perfect is that of who the active players are, and what options they have available to them. As we have modeled the game, only the firms are active players. This does not make much difference to the equilibrium analysis, since players are assumed to maximize their expected utility and thus are playing a best response if they always expect firms to offer the same policies. Indeed, even if firms are expected to go out of business in response to losses, it is never in a consumer's interest to behave otherwise. However, whether the consumers are active players does make a difference to the behavior of firms, which will be different if they anticipate that consumers will also make mistakes.

For example, it is easy to see that any separating equilibrium (which is the only sort that really permits consumer mistakes) can be upset by such mistakes, regardless of whether or not it involves

subsidization. The reason is that, while it is a large mistake for low-risk agents to purchase the policy intended for high-risk agents, the latter are indifferent between the policies they are supposed to buy and those offered to the low-risk agents. Thus, any firm that offers such a pair of policies will face negative expected profits if the consumers make small mistakes. This might lead to the choice of policies which earn a small positive profit if consumers do not make mistakes, but this is incompatible with equilibrium.

This is not really fatal, since it leads to the use of an ε-equilibrium concept that produces nearly the same results as before, but it does complicate matters.

For the balance of this section, we shall concentrate on the case of two firms who are the only active players in a market containing two types of customers.

First, let us consider the single-policy model. We take the strategy space to be the set of all nonnegative pairs dominated by (b_0,b_0). If two firms offer policies p and p', we can describe the profits of each as a function of both policies, as in Fig. 11.

If the rival firm locates in region A(p), the firm offering p will earn the same profits as previously. If the rival locates in C(p), the firm offering p will earn no profits. If the rival locates in B(p)

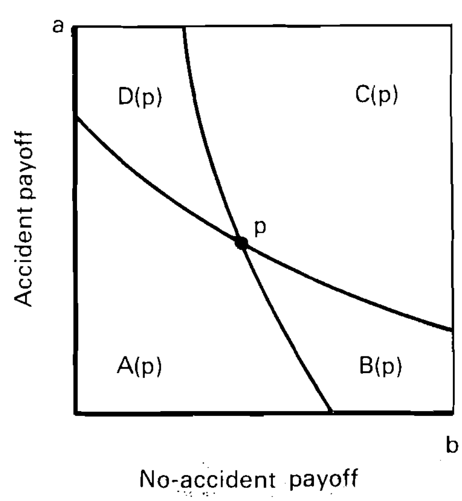

Fig. 11 -- Demands with two policies

(D(p)), the firm offering p will earn the profits (losses) it formerly earned on its low-risk (high-risk) customers alone. In addition to these data, we need to know the position of the breakeven lines to evaluate each of these possibilities.

To facilitate the discussion, let us denote the profits earned on the high- and low-risk customers by H1(p) and H2(p), where p = (a,b):

$$H1(p) = n[t_1(a_0 - a) + (1 - t_1)(b_0 - b)]$$

$$H2(p) = (1 - n)[t_2(a_0 - a) + (1 - t_2)(b_0 - b)]$$

We also need to classify the strategy space according to the profits earned on each type, as in Fig. 12.

Now we can show that the policy of full fair insurance for the high-risk types is undominated. Let this policy be p~ = (c(R),c(R)) (defined in Sec. II, Eq. (9)), and let p be any alternative policy the firm is considering.

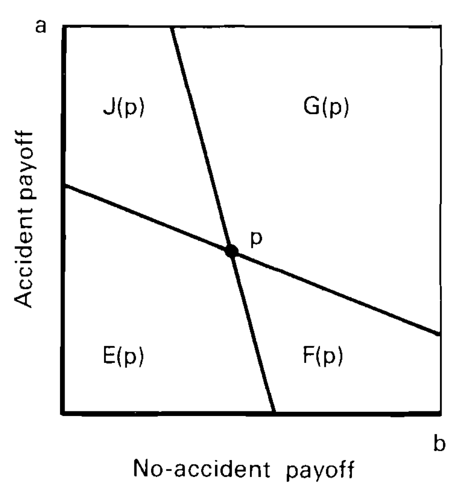

Fig. 12 -- Profits with two policies

If p is in F(p~) then $H_1(p) > H_1(p\sim)$ and $H_2(p) < H_2(p\sim)$
If p is in G(p~) then $H_1(p) < H_1(p\sim)$ and $H_2(p) < H_2(p\sim)$
If p is in J(p~) then $H_1(p) < H_1(p\sim)$ and $H_2(p) > H_2(p\sim)$

To show that p is undominated, it suffices to demonstrate that for each position p could occupy, there is a place q for the other firm to locate such that the resulting profits to the firm offering p~ strictly exceed those it could earn offering p.

For example, if p is in E(p~), then if q belongs to A(p~) and C(p), the profits if the firm offers p~ will be $H_1(p\sim) + H_2(p\sim) > 0$, while the profits if the firm offers p will be 0. This will always be possible since E(p~) belongs to A(p~).

In general, it is useful to classify the possible demands that could result from entry as follows.

If p is in:	then possible demands for p~ and p (respectively) are:
A(p~)	(12,1);(12,2);(12,0);(12,12);(0,0);(1,0);(1,1);(2,0);(2,2)
B(p~)	(12,12);(12,2);(0,0);(0,2);(1,0);(1,1);(1,2);(1,12);(2,2)
C(p~)	(12,12);(1,12);(1,1);(2,12);(2,2);(0,12);(0,1);(0,2);(0,0)
D(p~)	(12,12);(12,1);(0,0);(0,1);(1,1);(2,0);(2,1);(2,2);(2,12)

where 1 and 2 stand for types 1 and 2, respectively; 0 stands for the empty set; and 12 stands for both types 1 and 2.

There are eight possible combinations of utility (A-D) and relative profitability (E-J). They will be denoted by both letters; the intersection of A(p~) and E(p~) will be denoted AE(p~), etc.

If p is in:	then demands favoring p that are feasible are:
AE(p~)	(12,0); (2,0)
AF(p~)	(12,2); (12,0); (2,0); (2,2)
AJ(p~)	(12,0); (12,1); (1,1); (2,0)
BG(p~)	(12,12); (12,2); (2,2); (1,1)
BF(p~)	(12,2); (2,2)
CG(p~)	(12,12); (1,1); (2,2)
DG(p~)	(12,12); (12,2); (1,1); (2,0); (2,2)
DJ(p~)	(12,1); (1,1); (2,0)

Since none of these is empty, it is evident that p~ is undominated. The same cannot be said for (a(R),b(R)), the policy that is offered to the low-risk agents in the RS separating equilibrium. We shall use the technique developed above to determine the entire set of undominated single policies.

First, we may divide the policy space into profitability regions relative to the two classes of consumers; Fig. 13 shows this division.

There are nine regions, which may be described by the sign of the profits they would yield if policies in them were purchased by the type 1 and type 2 agents, respectively. Thus, (a_0,b_0) is in region (0,0); (c(R),c(R)) is in region (0,+); (a(R),b(R)) is in region (-,0); and both policies of a nonperverse subsidy equilibrium belong to region (-,+).

Next, for any candidate policy p~, we must classify the possible positions that a policy which dominates it could occupy. The two important characteristics are the relative profitabilities of the two policies for each group (determined by the categorization E-J), and the relative attractiveness of each policy (determined by the categorization

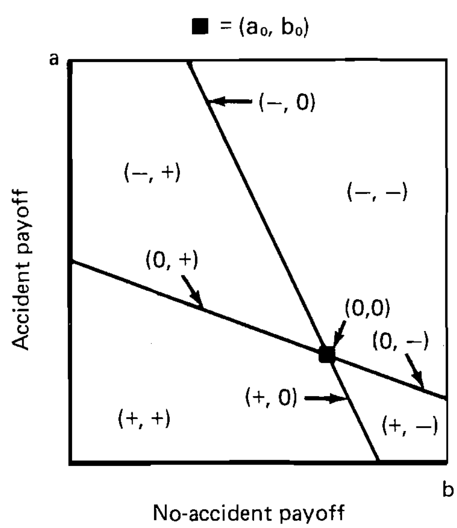

Fig. 13 -- Profits with two types

A-D). The latter is important even if the two policies never compete in the market, since it determines the differential allocation of demand in the face of entry.

The possibly nonempty regions and the relative profitabilities of p~ and p located in the indicated region relative to p~ are:

Region	$\max[H_1(p), H_1(p\sim)]$	$\max[H_2(p), H_2(p\sim)]$
AE	$H_1(p)$	$H_2(p)$
AF	$H_1(p)$	$H_2(p\sim)$
AG	$H_1(p\sim)$	$H_2(p\sim)$
AJ	$H_1(p\sim)$	$H_2(p)$
BF	$H_1(p)$	$H_2(p\sim)$
BG	$H_1(p\sim)$	$H_2(p\sim)$
CE	$H_1(p)$	$H_2(p)$
CG	$H_1(p\sim)$	$H_2(p\sim)$
CJ	$H_1(p\sim)$	$H_2(p)$
DE	$H_1(p)$	$H_2(p)$
DG	$H_1(p\sim)$	$H_2(p\sim)$
DJ	$H_1(p\sim)$	$H_2(p)$

This information can be combined with information as to the location of p~ to determine whether or not there is a q such that p~ can give a strictly higher payoff against q than p can.

Suppose, for example, that p~ was in region (-,-), and that p was in AE(p~). We would then have $H_1(p) > H_1(p\sim) < 0$ and $H_2(p) > H_2(p\sim) < 0$. In fact, it is easy to see that p can be chosen such that $H_1(p) > 0 > H_1(p\sim)$ and $H_2(p) > 0 > H_2(p\sim)$. It is clear that no matter where q is placed, p will afford at least equal profits to p~, and will often provide better profits. Thus p dominates p~. The same is true of p~ in (0,-), (-,0), and (0,0).

For the (-,+) region, if p is in E(p~), the profits earned with p~ will be higher if q is in D(p~) and in C(p). This intersection will be nonempty unless p itself belongs to C(p~), giving both types higher utility. Thus if CE(p~) is empty, no policy in E(p~) dominates p~. This will be the case, if the type 2 indifference curve through p~ is

steeper than bel. If p is in G(p~), then any q in E(p~) will belong to E(p), and will result in strictly larger profits for p~. If p is in F(p~), entry at q will result in larger profits for p~ if q is in D(p~) and either D(p) or C(p), both of which are nonempty. Finally, if p is in J(p~), then q will give p~ higher profits if it is in B(p) and not in A(p~) or if q belongs to both B(p) and C(p~). The former set is nonempty if a(p~) < b(p~); otherwise, the latter set is nonempty. Thus, policies in the (-,+) region are undominated by other single policies as long as CE(p~) is nonempty.

Similar arguments show that policies in the (+,0), (0,+), and (+,+) regions are undominated by single policies.

Now consider a policy p~ in the perverse subsidy region (+,-), and suppose that p belongs to E(p~); then p earns higher profits on both types than p~. The only pattern of demand that necessarily gives p~ higher profits is when p~ is sold to the low-risk types alone, while p is not purchased. Therefore, we would need to find q in both B(p~) and C(p); but these are disjoint sets. Hence no policy in this region is undominated.

The set of policies undominated by any single policy is thus the set of policies that earn positive profits if purchased by the low-risk agents, and which have the property that CE(p~) is nonempty. This latter condition can be interpreted geometrically, since it is equivalent to:

$$\frac{U'(b)}{U'(a)} \geq \frac{t_2(1-t_1)}{t_1(1-t_2)} \tag{1}$$

This is the equation of the ray through the origin defined by the condition that the low-risk indifference curves are tangent to the high-risk breakeven line. This condition is independent of n and of (a_0,b_0), so that several geometries are possible. A typical set of undominated policies is shown in Fig. 14.

It remains only to check that none of these policies is dominated by a mixed policy. First consider the policy p~ = (c(R),c(R)). Let p be any mixed policy that dominates p~. It must therefore be the case that p affords nonnegative profits when offered in competition with

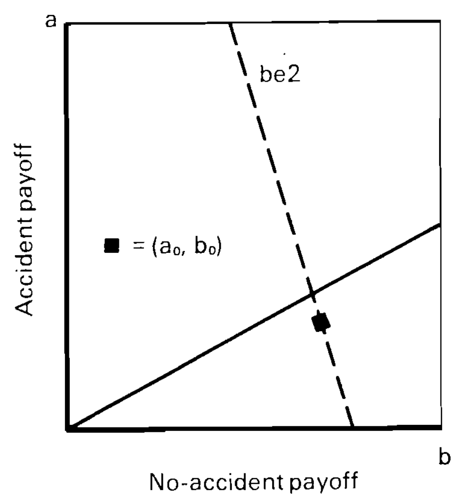

Fig. 14 -- A set of undominated policies

(a(R),b(R)). Since p~ is in equilibrium with this policy, everything in the support of p and therefore p itself must give exactly 0 profits, which in turn means that the support of p is contained in the set of policies not preferred to p~ by the high-risk types. Now suppose that firm 2 were offering a policy q slightly smaller in both coordinates than p~. By the argument used to show that p~ was undominated by any Pareto-inferior policy (the fact that it lies in region (0,+)), the support of p must be confined to the high-risk indifference curve running through p~ (and through (a(R),b(R))). Now consider any policy offered by the entrant that attracts no customers: Clearly, the incumbent's profits are strictly maximized by offering p~. Hence the support of p must be exactly p~.

In general, the payoffs $H_1(p)$ and $H_2(p)$ are strictly decreasing linear functions of p, and the payoff function H(p,q) can be defined as:

$$
(2)\ H(p,q) = \begin{cases} 0 & \text{if } p \text{ belongs to } A(q) \text{ [i.e., } q \text{ belongs to } C(p)] \\ H_1(p) & \text{if } p \text{ belongs to } D(q) \text{ [i.e., } q \text{ belongs to } B(p)] \\ H_2(p) & \text{if } p \text{ belongs to } B(q) \text{ [i.e., } q \text{ belongs to } D(p)] \\ H_1(p) + H_2(p) & \text{if } p \text{ belongs to } C(q) \text{ [i.e., } q \text{ belongs to } A(p)] \end{cases}
$$

Although this payoff function is not convex, it has a convex envelope as p varies among pure policies, so that if p can be dominated by a mixed policy, it can be dominated by a pure policy. Hence, the set of undominated single policies is equal to the set of policies that earn positive profits on at least one type.

The payoff function is different for Game II, which has anticipatory conjectures. For one thing, each of $H_1(p)$, $H_2(p)$, and $H_1(p) + H_2(p)$ is replaced by the maximum of itself and 0. Since exit is absorbing (no reentry is allowed), we can modify the payoff function according to the signs of $H_1(q)$, $H_2(q)$, and $H_1(q) + H_2(q)$. Denoting the sign-triple pattern as the triple $(\mathrm{sgn}(H_1(q)), \mathrm{sgn}(H_2(q)), \mathrm{sgn}(H_1(q) + H_2(q)))$, we have:

I. (+,+,+), (+,0,+), (0,+,+), or (0,0,0):

$$H(p,q) = \begin{cases} 0 & \text{if } p \text{ is in } C(q) \\ \max\,[H_1(p),0] & \text{if } p \text{ is in } D(q) \\ \max\,[H_2(p),0] & \text{if } p \text{ is in } B(q) \\ \max\,[H_1(p) + H_2(p),0] & \text{if } p \text{ is in } A(q) \end{cases}$$

II. (-,+,+), or (-,+,0):

$$H(p,q) = \begin{cases} 0 & \text{if } p \text{ is in } C(q) \\ \max\,[H_1(p),0] & \text{if } p \text{ is in } D(q) \\ \max\,[H_1(p) + H_2(p),0] & \text{if } p \text{ is in } A(q) \text{ or } B(q) \end{cases}$$

III. (-,+,-):

$$H(p,q) = \begin{cases} \max\,[H_1(p),0] & \text{if } p \text{ is in } D(q) \\ \max\,[H_1(p) + H_2(p),0] & \text{otherwise} \end{cases}$$

IV. (+,-,+), (+,-,0):

$$H(p,q) = \begin{cases} 0 & \text{if } p \text{ is in } C(q) \\ \max\,[H_2(p),0] & \text{if } p \text{ is in } B(q) \\ \max\,[H_1(p),0] & \text{if } p \text{ is in } A(q) \text{ or } D(q) \end{cases}$$

V. $(+,-,-)$, $(0,-,-)$:

$$H(p,q) = \begin{cases} \max\,[H_2(p),0] & \text{if } p \text{ is in } B(q) \\ \max\,[H_1(p) + H_2(p),0] & \text{otherwise} \end{cases}$$

VI. $(-,-,-)$:

$$H(p,q) = \max\,[H_1(p) + H_2(p),0]$$

The analysis of perfectness must take account of this variation in the payoff function. However, for the dominated regions of the game with Nash conjectures, q could be chosen anywhere, so that (c(R),c(R)) remains undominated and (a(R),a(R)) remains dominated.

We shall now see that the Wilson pooling equilibrium is undominated. Since it lies in the "normal subsidy" region, and earns zero profits overall, we may classify it as $(-,+,0)$ according to the above scheme. What we need to show is that it satisfies condition (I). However, at the pooling equilibrium the slope of the low-risk indifference curve equals that of the market breakeven line, which is steeper than the high-risk breakeven line.

The analysis of perfectness is more complicated for the multiple-option games. To simplify matters, we shall limit our attention to showing that policy pairs arising as interior subsidy equilibria are undominated.

The reason the analysis is different is that in the multiple-option model the industry structure is not well specified. In principle, if there are two firms in the market, one would expect both to offer the same pair of policies. Division of the profits between the two is irrelevant, since any equiproportional allocation gives each firm zero profits. However, if we assume that the other firm offers the same pair of policies, we can show that the set under consideration is undominated. Any other pair of policies that would make nonnegative profits if it served the whole market would attract no low-risk customers, since their indifference curve is tangent to the zero-profit locus. Hence, such a policy would make losses. On the other hand, a policy that would make positive profits if offered in competition with

the existing policy will necessarily inflict losses on the other firm, leading to exit and to negative profits.

In a future paper, we hope to examine the issue of perfectness in more depth.

V. EQUILIBRIUM AND OPTIMALITY IN A MANY-TYPE MODEL

In the two-type model, there were various ways of stating the reactive conjectures or strategies of the firm, all of which lead to the conclusion that an entrant whose entry left the incumbent firm with losses would find itself serving the whole population. This could be achieved by specifying that incumbent firms would:

- Drop all their policies in the face of any entry that resulted in losses;
- Drop all unprofitable policies in the face of loss-inflicting entry;
- Drop all unprofitable policies offered to agents who are higher-risk than the highest-risk type attracted by the loss-inflicting entrant; and so on.

These conjectures have different consequences in a model with more than two types, and in this section we examine the influence of various strategic or conjectural formulations.

We are using the words "strategic" and "conjectural" interchangeably here, since the reaction of incumbent firms to entry may be regarded either as a conjectured response or as a part of the firm's strategy that does not get used in equilibrium. The latter view is preferable from a game-theoretic point of view, since it removes some of the ad hoc nature of conjectures. We shall describe the market for insurance in a way that makes the strategic description a natural part of the model.

Essentially, we shall assume that a strategy for a firm consists of a portfolio of policies to be offered, together with a stipulation as to which of them will continue to be offered in the face of loss-inflicting entry. The reactive part of the strategy will depend on the entrant's policies, so that in a weak sense "the punishment" can be made to fit "the crime."

However, in keeping with the previous models, we shall not allow firms to change the terms of their existing policies. As we pointed out in the two-type model, allowing such freedom renders the equilibrium essentially indeterminate. For example, firms could offer any policies they wished (including those guaranteeing positive profits), backed by the threat of reverting to the subsidy equilibrium policies in the face of entry. Such a threat would certainly preclude entry, and would be perfectly credible. One might imagine a legal structure in which alteration of existing policies was forbidden by law.

In addition, we require that a potential entrant consider only the profits earned after incumbent firms complete their reactions. This is somewhat unrealistic, since it removes any of the "raiding" possibilities that one would expect if firms discount the future; essentially, one can think of this as forbidding such discounting, or as implying that incumbent reactions are instantaneous. If we allow discounting, the existence of equilibrium may be threatened. As firms become more myopic, the prospects an entrant faces more nearly resemble those faced when existing firms are constrained not to drop any policies. As we saw in the two-type model, this exacerbates the single-policy-per-firm nonexistence problem.

The organization of this section is as follows. First, we address the question of optimality in a many-type world. A simple geometric proof demonstrates that the maximum of any linear welfare function subject to separation and zero-profit constraints is attained at a unique portfolio which is a continuous function of the underlying parameters. Then we turn to the game described above. After describing the general model, we point out the special cases corresponding to the conjectural assumptions embodied in Rothschild-Stiglitz, Spence, and the straightforward generalization of the assumption we made in Sec. III. Then we define various types of equilibria for the general game. Nash equilibrium is obtained by requiring that, for every entrant there is a set of types containing those attracted by the entrant such that the entrant makes nonpositive profits on that set. A Credible Nash equilibrium is defined as a situation where to every loss-inflicting entrant there corresponds a set of types containing those attracted by

the entrant, such that the entrant earns nonpositive profits on that set and such that the incumbent portfolio earns nonnegative profits on the remaining types. Finally, a Trembling-hand Perfect Nash Equilibrium is defined to be an equilibrium robust against small mistakes on the part of the players. For the balance of the section, we focus on a postulated reaction in which all policies are dropped in response to loss-inflicting entry, and no policies are dropped if the entrant does not inflict losses. If the equilibrium policies are offered by a single firm, we can easily turn statements about the profitability of potential entry to statements about the feasibility of portfolios that improve the utility of various groups. We therefore investigate the optimality properties of the equilibria relative to these reactions, which are necessarily credible and, as we shall show, trembling-hand perfect.

The model consists of k "types" of agents. There is a proportion n(i) of players of type i, each of whom has a probability t(i) of having an accident, where:

$$t(1) > t(2) > \ldots > t(k)$$

Each has the same utility-for-certain-income function U, which is a strictly increasing, strictly concave, continuous function from the nonnegative real numbers to the real numbers. The expected utility function of type i for a state-contingent bundle (a,b) is:

$$U_i(a,b) = t(i)U(a) + (1 - t(i))U(b) \tag{1}$$

Optimality in the k-type world can be defined similarly to optimality in the two-type world. We shall not repeat the definition of E1 optimality, but we shall restate the E2 and E3 definitions, and add another constrained optimality problem that provides a sharper characterization of the RS separating contracts:

$$\max \Sigma_i \, w(i)Ui(a(i),b(i)) \text{ subject to:}$$

$$\text{(E2')} \quad \text{(i) } \Sigma_i \, n(i)[t(i)(a(i) - a(0)) + (1 - t(i))(b(i) - b(0))] \le 0$$

(ii) for each i, and for each j in $\{0, \ldots, k\}$

$$U_i(a(i),b(i)) \geq U_i(a(j),b(j))$$

This problem maximizes social welfare subject to an overall zero-profit constraint (which may be regarded as a feasibility condition) and subject to separation of types (which may be regarded as consumer sovereignty).

$$\max \Sigma_i \; w(i)U_i(a(i),b(i)) \text{ subject to:}$$

(E3') (i) $\Sigma_i n(i)[t(i)(a(i) - a(0)) + (1 - t(i))(b(i) - b(0))] \leq 0$

(ii) for each i, and for each j in $\{0, \ldots, k\}$

$$U_i(a(i),b(i)) \geq U_i(a(j),b(j))$$

(iii) $U_i(a(1),b(1)) \geq U(t(1),a(0)) + (1 - t(1),b(0))$

This problem bears the additional constraint that the highest-risk types get at least as much as they would under full fair insurance. An alternative optimization problem that captures the effect of competition is obtained by replacing the overall breakeven constraint (i) with the constraint that no type generates losses.

$$\max \Sigma_i \; w(i)U_i(a(i),b(i)) \text{ subject to:}$$

(E4') (i) for each i $t(i)(a(i) - a(0)) + (1 - t(i))(b(i) - b(0)) \leq 0$

(ii) for each i, and for each j in $\{0, \ldots, k\}$

$$U_i(a(i),b(i)) \geq U_i(a(j),b(j))$$

For the two-type case, problems E3' and E4' are equivalent.

Just as in the two-type case, the Ei optima are single-valued continuous functions of the underlying parameters. We shall give a different proof, which is more elegant and does not require differentiability of U. We begin by illustrating the proof for the two-type case.

First, recall the definition of the locus of feasible points for problem E2. These were shown in Fig. 5, which is reproduced below. It is not convex, but it can be shown to be relatively convex to the low-risk types' indifference curves.

Since the high-risk types always obtain full insurance by Proposition 2.1, we can represent this locus directly as a set of (a,b,c) triples. The zero-profit constraint (i) above is a simplex, while the separation constraint (ii) is everything above the "hill" which is obtained by lifting the high-risk indifference curves from the (a,b) plane until each is at the height c, such that all (a,b) pairs give utility U(c). The feasible region is the very non-convex intersection of these two shaded in Fig. 15.

In order to obtain the convexity result relative to the low-risk types' indifference curves, we shall transform this diagram by applying the utility function to each axis independently. The triple (a,b,c) thus becomes the triple (U(a),U(b),U(c)). The simplex representing the breakeven constraint (i) becomes a strictly convex "Pareto" set, while the "hill" representing the separation constraint (ii), now becomes a

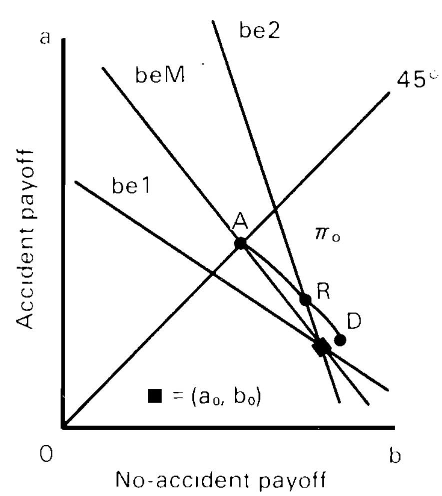

Fig. 5 -- The zero-profit locus

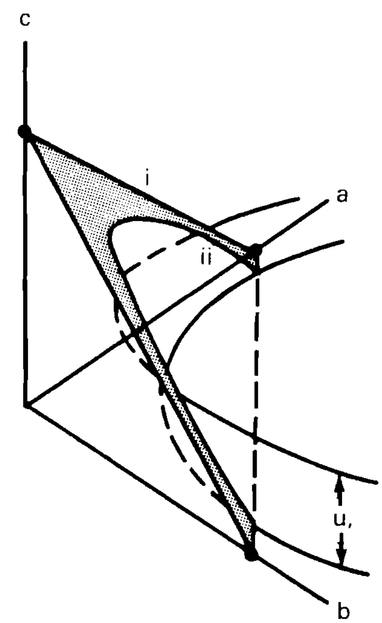

Fig. 15 -- Feasibility in policy space

plane constraint. It is quite clear that the feasible set is now a convex compact Pareto surface. This is the shaded region in Fig. 16.

The objective function is a hyperplane, and the uniqueness and continuity of the optimum follow from this fact. Figure 17 shows the position of the optimum for one special case, corresponding to

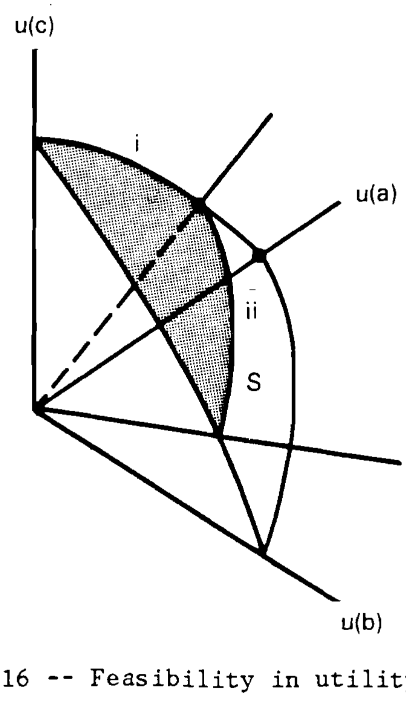

Fig. 16 -- Feasibility in utility space

maximizing only the low-risk types' utility, and for an "interior" case where social welfare includes the utility of the high-risk types.

Figure 18 illustrates the effect of changing the proportion of the two types. The breakeven constraint is the only one affected, and it "rotates" through the curve defined by actuarial fairness (indicated F). The area above this line is the area of normal subsidies, while the area below the line is occupied by perverse-subsidy policies. This picture makes it clear that in general, an interior optimum with normal subsidies will approach curve F as the proportion of high-risk agents rises, and will pass through as it increases further. For optima corresponding to problem E3, the feasible set is truncated at curve F, and the optimum will therefore remain there (for $n > n^*$) as a corner solution.

Now, let us focus our attention on the special case where the objective function only depends on the utility of the low(est)-risk types (see Fig. 19). In this case the optimum will lie on the curve S, which is the intersection of the separation plane and the breakeven Pareto set, and which forms the lower boundary of the feasible set. S always intersects F, and always from the same direction. For problem E3, it is the points lying above both F and S that are feasible.

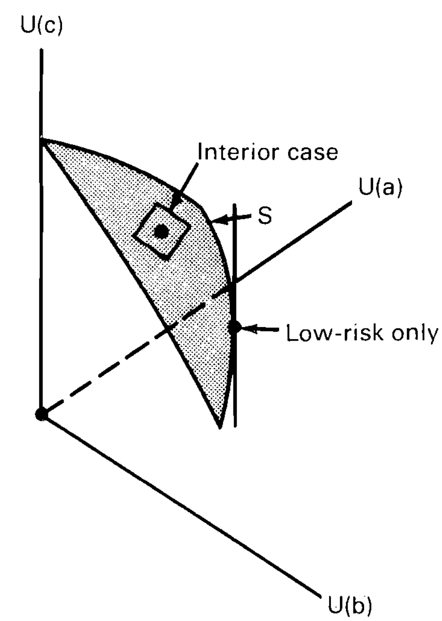

Fig. 17 -- Optimality in utility space

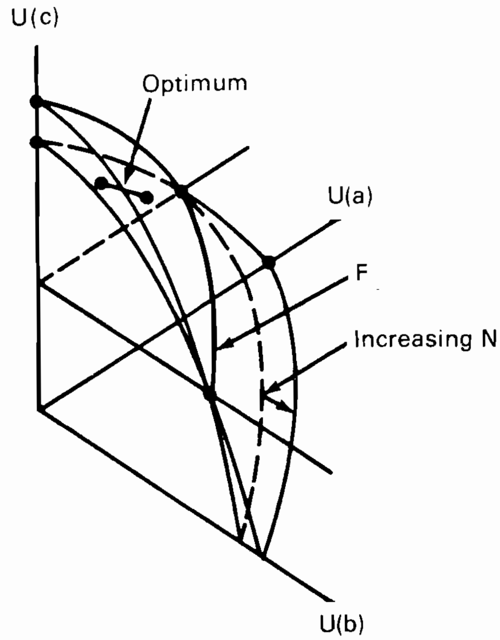

Fig. 18 -- The effect of population changes

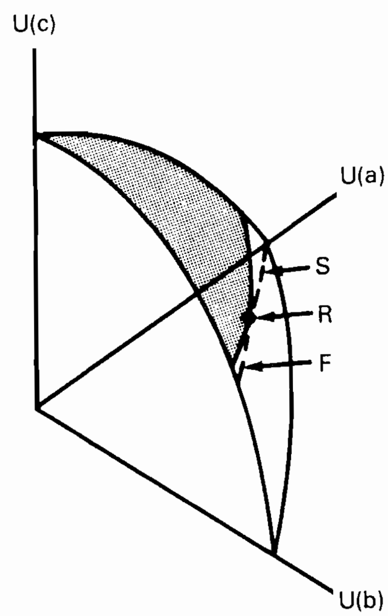

Fig. 19 -- E3 feasibility

As n(1) increases, curve F will remain constant, but curve S will shift as well, shifting down where it lies above F and up where it lies below F. It is easiest to visualize this in the (U(a),U(b)) plane. Figure 20 shows how the projection of the lower boundary of the feasible set shifts with increases in the proportion of high-risk agents. This diagram makes it clear why the E3 optimum moves smoothly along the shifting S curve for sufficiently low values of n(1), but remains stuck at the intersection point (the RS separating contracts) for higher n(1).

The formal arguments corresponding to this diagrammatic exposition follow. We need to show that the usual property that a linear function has a unique maximum on a strictly convex set continues to hold, if we impose additional linear constraints that do not affect the strict convexity of the relevant part of the set. Consider the maximization of a linear function on the 3-sphere. If we impose a constraint such as:

(2) $$x(1) \leq x$$

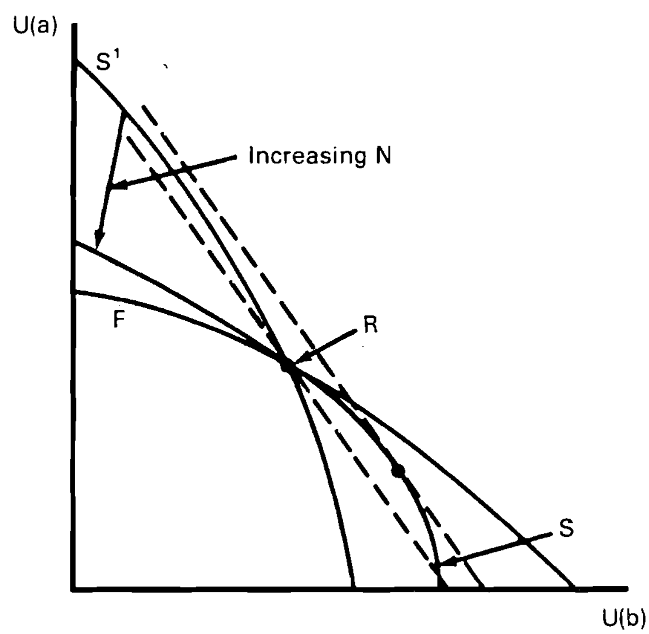

Fig. 20 -- E3 corner optimality

the resulting set is not strictly convex. However, the Pareto set of the set of feasible points has the property that nontrivial linear combinations of its members are not in the Pareto set, and the solution always belongs to the Pareto set; hence the solution is unique. This situation is illustrated in Fig. 21.

For any set H, let us define the strict Pareto set of H, PO(H), as:

(3) $\quad PO(H) = \{x \text{ in } H\text{: for all } y \text{ in } H\text{, there is } i \text{ s.t. } x(i) > y(i)\}$

The weak Pareto set of H, P~(H), is defined by:

(4) $\quad P\sim(H) = \{x \text{ in } H\text{: for all } y \text{ in } H\text{, there is } i \text{ s.t. } x(i) \geq y(i)\}$

The result we need is:

Proposition 5.1: Consider the maximization problem

max wx subject to: x in C,

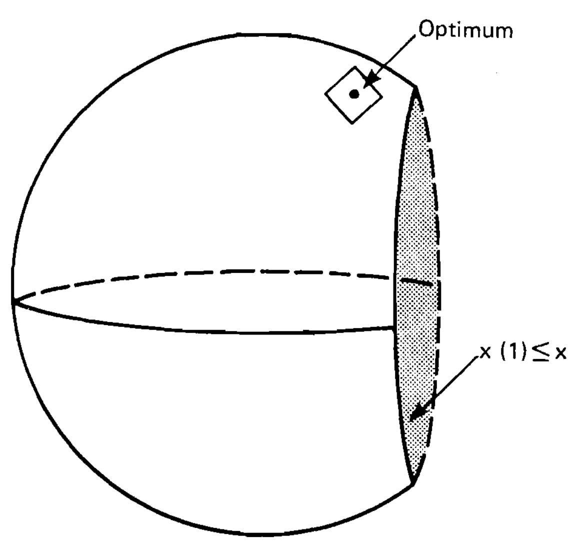

Fig. 21 -- Relative strict convexity

where $w(i) \geq 0$, all i, C is a convex compact set, and P~(C) has the strict convexity property that x,y in P~(C), s in (0,1) implies:

$$sx + (1 - s)y \text{ not in } P\sim(C)$$

Then there is a unique solution to the problem.

PROOF: Since the w(i) are all nonnegative, the solution belongs to P~(C). If there were two solutions, x and y, we would have:

$$wx = wy = w[sx + (1 - s)y]$$

for all s in [0,1]. Therefore, since C is convex, $sx + (1 - s)y$ must also be a solution. By the strict convexity property of P~(C), it does not belong to P~(C) for s in (0,1). This is a contradiction. QED

Using this result, we can show that there are unique solutions to the various optimization problems posed above.

Proposition 5.2: The optimization problem E2 has a unique solution, which is a continuous function of the N(i) and t(i) as long as the ranking

(5) $$t(1) > t(2) > \ldots > t(k)$$

is preserved.

PROOF: The set of points satisfying constraint (i) is a convex polyhedron. By strict concavity of the utility function U, the image of this set under the map (U, ..., U) [2k times] is a strictly convex compact set. If the vector of utilities in the generic point is written

$$(u,v) = [u(1),v(1),u(2), \ldots, u(k),v(k)]$$

where $u(i) = U(a(i))$ and $v(i) = U(b(i))$, and if the vectors of proportions n(i) and accident probabilities t(i) are written N and T,

respectively, this set can be written in the form G(N,T), for some continuous, strictly convex-valued correspondence G. The image of the set of points satisfying the separation constraint (ii) under (U, ..., U) can be written F(T), for some polyhedron-valued correspondence F continuous on the set of T satisfying (3) above. Indeed, if we write G(N,T) as

$$(6) \qquad g(u,v,N,T) \leq 0,$$

it is easy to see that g is a continuous function, strictly increasing and strictly convex in u,v.

In light of the ranking (3), the equations defining F(T) are of the form:

$$(7) \qquad t(i)u(i) + (1 - t(i))v(i) \geq t(i)u(i + 1) + (1 - t(i))v(i + 1)$$

for i = 1, ..., k-1. If we denote the intersection of G(N,T) and F(T) as M(N,T), we can see that P~(L(N,T)) has the strict convexity property referred to in Proposition 5.1. The reason is that the inverse of the utility function U is a strictly convex function w. Therefore, if (u,v) and (u',v') satisfy Eq. (7), and if s is in (0,1) then:

$$(8) \; \Sigma_i \, n(i)[t(i)w(su(i) + (1 - s)u'(i)) + (1 - t(i))w(sv(i) + (1 - s)v'(i))]$$

$$< t(i)a(0) + (1 - t(i))b(0)$$

so that s(u,v) + (1 - s)(u',v') is not in P~(G(N,T)).

Of course, there is a further set of constraints to consider--the individual rationality constraints:

$$(9) \qquad Ui(a(i),b(i)) \geq Ui(a(0),b(0));$$

but these will not affect the result as long as all types have nonnegative weight in the objective function.

Moreover, since the feasible set is a continuous correspondence of N and T, the value of the maximum welfare is a continuous function of N and T, and the optimal policy correspondence is uppersemicontinuous in N and T. Since it is also single-valued, it is continuous.
QED

Examination of the additional constraints (iii) and (iv) shows that they do not affect the strict convexity properties of the Pareto-optimal feasible set, leading to:

Corollary 5.3: The E3 and E4 optima are unique and continuous in N and T.

At this point, it is useful to contrast the notion of optimality used here with that used by Spence. He defines a sequence of optimization problems, which may be expressed in terms of the present model as follows:

(F1) $q(1) = U(t(1)a(0) + t(2)b(0))$

(Fi) $q(i) = \max t(i)U(a(i)) + (1 - t(i))U(b(i))$ subject to

$t(j)U(a(j)) + (1 - t(j))U(b(j)) \geq t(j)U(a(j + 1)) + (1 - t(j))U(b(j + 1))$ for all $j < i$

$t(j)U(a(j)) + (1 - t(j))U(b(j)) \geq q(j)$ for all $j < i$

$n(j)[t(j)(a(j) - a(0)) + (1 - t(j))(b(j) - b(0))] = 0$

and focuses on optimality in the Fk sense. By adapting the technique of proof used above, we can obtain:

Proposition 5.4: There is a unique Fk-optimal set of policies, and it is a continuous function of the parameters N and T.

PROOF: The utility constraints given by the $q(j)$ $j < i$ for each problem Fi do not affect the nonemptiness or the strict convexity of the weak Pareto set in each problem Fi. Thus each Fi has a unique continuous solution as a function of N, T and the $q(j)$. However, the $q(j)$ are also continuous functions of N and T. QED

Now we turn to the strategic formulation. First, let us denote a generic portfolio of contracts by $w = w(1), \ldots, w(k)$, where for each i, $w(i) = a(i),b(i)$ and for all i,j:

$$(10) \qquad U(i,w) = U(i,w(i)) \geq \max\{U(i,w(j)), U(i,(a(0),b(0))\}$$

This is quite general, since any portfolio can be represented as such a "separating portfolio" by adding redundant contracts if there is pooling, and removing unwanted contracts. The set of such portfolios is W. We can also define the profit that is earned on type i, when w is offered, to be:

$$(11) \qquad P(i,w) = n(i)[t(i)(a(0) - a(i)) + (1 - t(i))(b(0) - b(i))]$$

For any subset S of $\{1, \ldots, k\} = K$, we define

$$(12) \qquad P(S,w) = \sum_{i \text{ in } S} P(i,w)$$

If there is a set F of active firms, each of which offers (wlog) a separating portfolio $w(\ :f)$, consumer choice will determine an effective portfolio $w(\ :F)$. Of course, this may contain more than one policy for each consumer, but at equilibrium this will not be true; only that policy offering the most profit on each type will be offered. In addition, at equilibrium the profits earned by each firm will be zero, so we shall find it convenient for most of this section to assume that all firms offer the same portfolio $w(\ :F)$, which we shall refer to as "the" incumbent portfolio.

If a portfolio w is being offered in a market and an entrant offers a new portfolio w', we can divide the types into those who will switch to the new policy $w'(i)$ and those who continue to purchase $w(i)$. Let us

denote by $C(w';w)$ the clientele of the entrant offering w'. In keeping with the tradition of Nash equilibrium, we shall assume that consumers switch policies only when they strictly prefer the new policy to the old.

$$C(w';w) = \{i: U(i,w') > U(i,w)\} \tag{13}$$

A strategy for a firm is a pair consisting of a portfolio w and a reaction which is a map R from the space of portfolios W to the set of subsets of K satisfying

$$\begin{aligned} &\text{(i) } R(w') \text{ is a subset of } K - C(w';w); \text{ and} \\ &\text{(ii) } R(w') = K - C(w';w) \text{ if } P(K - C(w';w),w) \geq 0 \end{aligned} \tag{14}$$

The interpretation is that the firm will offer those policies corresponding to types in $R(w')$ in the face of an entrant offering w'. The first condition allows us to simplify the incumbent's post-entry profits to $P(R(w'),w)$, while the second condition ensures that only loss-inflicting entry will draw a reaction. The second condition also means that the assumption that all firms active in equilibrium can be treated as offering the same contract is consistent with reaction, since in that case $P(K - C(w';w),w) = P(K,w) = 0$.

Definition 5.5: (w,R) is a Nash equilibrium iff $P(K,w) = 0$, and for every w',

$$P(K - R(w'),w') \leq 0. \tag{16}$$

(w,R) is a Credible-Reaction Nash Equilibrium iff $P(K,w) = 0$ and, for every w'

$$\begin{aligned} &P(K - R(w'),w') \leq 0; \text{ and} \\ &P(R(w'),w) \geq 0. \end{aligned} \tag{17}$$

We shall find it convenient to suppress R when we discuss optimality, so we shall rephrase the above definition in terms of the portfolio w alone.

Definition 5.6: w is a Nash portfolio iff $P(K,w) = 0$ and for every w' there exists a subset $R(w')$ of $K - C(w';w)$ such that

$$P(K - R(w'),w') \leq 0. \tag{18}$$

w is a credible Nash portfolio iff $P(K,w) = 0$ and for every w' there exists a subset $R(w')$ of $K - C(w';w)$ such that

$$\begin{aligned} &P(K - R(w'),w') \leq 0; \text{ and} \\ &P(R(w'),w) \leq 0. \end{aligned} \tag{19}$$

We have not made a distinction between loss-inflicting entry and other entry. The reason is that a portfolio w offering positive profits can be Pareto-dominated by another positive-profit portfolio w'. In this case, $C(w';w) = K$, and w cannot be a Nash portfolio. If w offers zero profits, then the existence of a w' that does not inflict losses and makes nonnegative profits means that w can be Pareto-dominated by a positive-profit portfolio; essentially, this says that w was not profit-maximizing on its unprofitable customers given the separation constraint. In either case, the existence of a w' which does not inflict losses means that there exists a w' which captures the entire market and makes positive profits.

There remains the question of trembling-hand perfection. Fortunately, it is possible to find trembling-hand perfect equilibria without introducing mixed strategies explicitly. In games with a finite number of pure strategies, a trembling-hand perfect equilibrium is defined to be the limit of a sequence of ε-perfect equilibria as ε tends to 0. An ε-perfect equilibrium is an n-tuple of completely mixed strategies with the property that no pure strategy not belonging to the best-response correspondence of a player is used with probability exceeding ε. For two-player games, the definition of a trembling-hand perfect equilibrium is entirely equivalent to the condition that neither player uses a weakly dominated strategy with positive probability. In games with more than two players, the "undominated-strategy" condition is necessary but not sufficient. The reason is that the particular

combination of the other player's strategies against which an undominated strategy is the unique best reply may not be attainable when the other players cannot correlate their strategy choices. Under these circumstances, any sequence of ε-perfect equilibria may eliminate the use of a particular strategy even when it is undominated. However, in a wide class of games including the present one, each player's payoff depends on the other players' strategies only through an aggregate of those strategies. In other words, what all players can do by correlation, a single player could do by a mixed strategy. This in turn implies that the conditions of trembling-hand perfection and of equilibrium in undominated strategies are equivalent.

In games with a continuum of pure strategies, the notion of a completely mixed strategy is complicated by the fact that a mixed strategy may be a combination of a density and some atoms of probability. If we represent such a strategy as a cumulative distribution function, we obtain a weakly increasing function with jumps in it. The payoffs at atoms of probability are not comparable with those at continuity points of the cumulative distribution, and we must reinterpret both "completely mixed" and "domination" in terms of these jumps. This is done in Cave (1984). Essentially, the idea is that ε-perfect equilibria converging to an equilibrium with jumps are required to have the same set of jumps as the limit point. Strategies that are continuity points are required to be used with positive density not exceeding ε if they are inferior to another continuity-point pure strategy, and jump-point strategies are required to be used with positive probability not exceeding ε if they are inferior to another jump-point strategy. On the other hand, domination can also be defined with respect to a set of jumps. For our purposes, it is sufficient to note that if this is done carefully, an equilibrium in pure strategies will be trembling-hand perfect iff it does not use any dominated strategy with positive probability.

Definition 5.7: (w,R) is a Trembling-hand Perfect Nash Equilibrium if it is a Nash equilibrium, and for every (w',R') there exists w~ such that

$$P(R(w\sim),w) > P(R'(w\sim),w') \quad (20)$$

We can associate certain special cases of the reactive part of a strategy with the equilibrium concepts proposed by other authors. For example, the Nash equilibrium in policies suggested by Rothschild-Stiglitz corresponds to:

$$R(w') = K - C(w';w) \tag{21}$$

A straightforward generalization of the reactive strategy used in Sec. III entails:

$$R(w') = \begin{cases} K - C(w';w) & \text{iff } P(K - C(w';w),w) \geq 0; \\ 0 & \text{otherwise.} \end{cases} \tag{22}$$

Finally, the reactive strategy suggested by Spence is:

$$R(w') = K - C(w';w) - \{i\colon P(i,w) < 0 \text{ and } i < \min\{j\colon j \text{ in } C(w';w)\}\} \tag{23}$$

Neither (21) nor (23) necessarily satisfies credibility, but (22) is credible by definition. In addition, there may well be no equilibrium relative to the R defined in (21), while there is always an equilibrium relative to the R defined in (22) or (23). This, incidentally, shows that there is always an equilibrium in the general game. We shall demonstrate existence of a credible equilibrium in the general game by demonstrating existence of an equilibrium with respect to (22).

Proposition 5.8: Let R be as defined in (22). Then there exists w such that (w,R) is a Credible-Reaction Nash Equilibrium.

PROOF: The proposition is equivalent to the assertion that there exists w such that both L(w) and M(w) ≤ 0, where

$$\begin{aligned} L(w) &= \max P(C(w';w),w') \text{ s.t. } P(C(w';w),w) \leq 0 \\ M(w) &= \max P(K,w') \text{ s.t. } P(C(w';w),w) > 0 \end{aligned} \tag{24}$$

It is not difficult to eliminate the condition that $L(w)$ be nonpositive, since $L(w) > 0$ is equivalent to the existence of a subset S of types on which the incumbent earns nonpositive profits, but on which the incumbent could earn positive profits while improving the welfare of every member of S and offering the same policies as before to types not in S. By continuity, it follows that there exists a policy w' such that $C(w';w) = K$ and $P(K,w') > 0$. Therefore, $L(w) > 0$ implies $M(w) > 0$.

Hence we may confine ourselves to showing that there exists w for which $M(w)$ is nonpositive. This is equivalent to the existence of a fixed point to the following correspondence:

$$T(w) = \arg\max P(K,w') \text{ s.t. } P(C^*(w';w),w) \geq 0 \quad \text{where}$$
$$\text{(25)} \qquad C^*(w';w) = \{i: U(i,w') \geq U(i,w)\}$$

since if w is such a fixed point, any w' attracting (even weakly) a profitable (even weakly) set of consumers cannot offer positive profits. $P(K,w')$ is a linear function of w', so we must turn our attention to the set of separating portfolios and to the function C. It is clear that $T(w)$ is u.s.c. and nonempty-valued, since it represents the maximum of $P(K,w')$ over

$$\text{(26)} \qquad D(w) = \{w': P(C^*(w';w),w) \geq 0\}$$

which is compact, and continuous in w. However, $D(w)$ will not always be a convex set, but this nonconvexity does not really matter. In the two-type case, the fact that the high-risk agents always receive full insurance allows us to limit our attention to a convex-valued subcorrespondence of D. The same device works in the many-type case.

For each subset S of K such that $P(S,w)$ is nonnegative, we can limit our attention to $h(S,w) = \arg\max P(S,w')$ s.t. $C^*(w';w) = S$; in other words, the maximum of $P(K,w')$ over $D(w)$ will agree with the maximum of $P(K,w)$ over

$$\text{(27)} \qquad D^*(w) = \{w' \text{ in } D(w): w' \text{ belongs to } h(C(w';w),w)\}$$

which is the union of finitely-many policies (by Proposition 5.1), each of which varies continuously with w. Therefore, T(w) satisfies the conditions of the Kakutani fixed-point theorem, and there exists a fixed point. QED

This proposition also serves to demonstrate existence of equilibrium for the general-reaction case, since policies that are in equilibrium with the reactions used above will be equilibrium policies in general. In the same way, this proposition demonstrates the existence of a credible-reaction equilibrium in general. In addition, the proof of the proposition shows that such an equilibrium is undominated; a dominating portfolio could be offered by an entrant. In this case, the requirement of trembling-hand perfectness does not add to the requirement of subgame-perfectness embodied in the credible-reaction condition.

The equilibrium statement that the incumbent must be maximizing profits on any subset of types for which his profits are nonnegative can be turned into a dual statement which characterizes the welfare properties of equilibrium.

In particular, we can associate to each subset (or coalition) S the set of utility allocations that could be achieved subject to separation and subject to a profit constraint. In effect, the equilibrium portfolio is a zero-profit portfolio with the requirements of separation and the requirement that there does not exist another separating portfolio that gives higher utility to every member of a coalition on which the equilibrium earns nonnegative profits, and which earns nonnegative profits overall. This is a core-like condition, and the equilibria of these games can all be regarded as satifying such conditions.

In the two-type case, the equilibrium under both Spence's and our reaction specification maximized the utility of the low-risk types subject to separation, zero profits, and giving the high-risk types as much as they could get with full fair insurance. Spence's generalized equilibrium possesses Fk optimality (see above), which is the same as the core when each type j has "value" V(j) = q(j). Coalitions of more than one type have a characteristic function given by the vector of q(j)'s.

Let us define $U(S,w)$ to be the vector $[U(i,w): i \text{ in } S]$, and $U(S)$ to be a vector of utilities, one for each member of S. We can associate two characteristic functions to a given coalition. The first is relative to a given initial allocation w:

$$(28)\ V(S,w) = \{U(S): \text{there exists } w' \text{ with } U(S) \geq U(S,w');\ P(S,w') \geq P(S,W)\}$$

The second is not related to any particular allocation:

$$(29)\ \ X(S) = \{U(S): \text{thère exists } w' \text{ with } U(S) \geq U(S,w');\ P(S,w') \geq 0\}$$

The second is the characteristic function appropriate to a world of complete information; it cannot be used directly here, since it requires that each type be given its full fair insurance utility, which is impossible. On the other hand, it is the appropriate benchmark to be applied to coalitions that are unprofitable under the equilibrium w. Indeed, the optimality property of equilibrium can be stated in functional form as a corollary to Proposition 5.1. First, let us write:

$$(30) \qquad E(S:w) = \{w': C(w';w) = S\}$$

for the set of portfolios that would attract S.

Proposition 5.9: w is an equilibrium portfolio relative to reactions (22) iff for every S such that $P(S,w) \geq 0$:

$$(31) \qquad E(S:w) \text{ is contained in } \{w': P(K,w') \leq 0\}$$

PROOF: This is essentially the half of the definition of equilibrium corresponding to entrants who inflict losses on the incumbent. It is therefore necessary. To see sufficiency, suppose that for some S with $P(S,w) \leq 0$, there existed w' in $E(S:w)$ such that $P(S;w') > 0$.
Now consider the portfolio w' defined by

$$w''(i) = \begin{cases} w(i) \text{ for } i \text{ in } K/S \\ w'(i) \text{ for } i \text{ in } S \end{cases}$$

This will be weakly preferred by all types and gives positive profits equal to $P(S,w') + P(K/S,w)$. By continuity, there exists another portfolio which gives positive profits and is strictly preferred by all types--a contradiction. QED

The highest-risk types can always be separated, so the equilibrium will give them at least their full fair insurance utility, which means that the equilibrium policy will lose money on the highest-risk types. If the equilibrium policy loses money on $S = \{1,2\}$, then 1 and 2 must get at least their two-type subsidy-equilibrium utilities. If the equilibrium policy makes money on $S = \{1,2\}$, an entrant could offer any pair of policies that are a marginal improvement over the "equilibrium" offering to S; this would attract S and induce exit of the incumbent policies, which would leave the entrant with positive profits overall.

This line of reasoning shows that for any set S of the form $S = \{1, \ldots, s\}$, it cannot be the case that $P(S,w^*) > 0$. If $P(S,w^*) < 0$, then each member of S must get at least the utility he gets in the equilibrium when the set of types is S. We can therefore state a sharpened form of the optimality problem solved by the equilibrium.

Proposition 5.10: Define a sequence of utility K-vectors, $r(i,j)$ by:

$$r(i;j) = U(i,w\sim(\ ,j)) \text{ if } i \leq j,$$

$$0 \text{ otherwise},$$

where

$$w\sim(j) \text{ solves } \max_{w} (U(j,w) \text{ subject to})$$

(i) w a separating portfolio;

(ii) $P(\{1, \ldots, j\},w) \geq 0$

(iii) $U(i,w) \geq r(i;j - 1)$ for all $i < j$

Then w^* is a Nash equilibrium relative to reactions (22) iff

$$U(i,w^*) \geq r(i,k) \text{ for all } i$$

where k is the lowest-risk type.

PROOF: First, suppose that w^* does not dominate $r(i,k)$. Then there is a coalition S of the form $\{1, \ldots, s\}$ and a portfolio $w\sim$ such that

$$U(i,w\sim) > U(i,w^*) \text{ for all } i \text{ in } S$$
$$P(S,w\sim) > 0$$
$$P(K - S,w^*) \geq 0$$

Not only will an entrant enter successfully with w' but such an entrant could construct a Pareto-dominating profitable profile w' and enter with that.

The reverse implication can be verified from the fact that there is a unique solution to the condition that $w\sim$ dominates $r(i,k)$. This follows from Proposition 5.1. If every equilibrium dominates $r(i,k)$, and if there exists an equilibrium, it must be the unique feasible policy that dominates $r(i,k)$. QED

We have therefore shown the uniqueness of the equilibrium relative to (22), as well as existence and optimality. The existence of equilibria relative to the general reaction strategies is thus assured. However, the uniqueness and the consequent equivalence of the optimality criteria to which they give rise remain an open question.

VI. DISCUSSION AND OPEN QUESTIONS

This report has described the implications of varying the strategies available to players in a simple class of games of asymmetric information, and has attempted to put existing results into a game-theoretic context where the implications of various strategic assumptions can be examined. Throughout, we have maintained the point of view that the different "equilibrium" concepts (Nash equilibrium, reactive equilibrium) can best be viewed as Nash equilibria of different games. Indeed, the reactive equilibria have an additional property that is analogous to subgame perfection. We have exploited these properties to determine the implications of allowing firms to offer more than one contract each, with the free-entry constraint operating only at the aggregate level. This allows firms to internalize the externality caused by asymmetries in information, and leads to constrained optimality. Of course, it might be argued that the insurance firms in Rothschild-Stiglitz, Wilson, Riley, et al., could offer more than one policy as long as each policy broke even in expectation. This is the reason why we stress the game-theoretic setting, since such behavior by firms is correctly described by Riley as "Walrasian" and may well be inappropriate to game-theoretic models.

We found several reasons to prefer an equilibrium with subsidization to the separating Walrasian contracts. In the first place, for firms offering a single policy each, the separating policy offered to the low-risk types is weakly dominated, and thus cannot form part of a perfect equilibrium. If a perfect equilibrium exists for the single-policy game, it is probably the mixed-strategy equilibrium found by Rosenthal and Weiss. That equilibrium has its support confined to the set of undominated policies and has implicit subsidization. By contrast, the subsidy equilibrium is perfect.

The second problem has to do with existence. If we stipulate that the strategies of the firms consist only of portfolios of policies, without any reaction, then Nash equilibrium may not exist. This was already shown for the single-option case by Rothschild and Stiglitz, who

show that for a range of population distributions, the Walrasian separating policies can be beaten by a "pooling" policy. The latter, however, is vulnerable to "skimming," and thus, cannot itself be an equilibrium. If we allow firms to offer more than one policy, this nonexistence problem is exacerbated, and Nash equilibrium in pure strategies is limited to the Walrasian separating contracts for cases where they are efficient.

The nonexistence problem can be alleviated by changing the strategy space, and uniqueness guaranteed by tightening the solution to a perfect Nash equilibrium. For the single-policy case, this was carried out by Wilson, who stipulated that equilibrium should be measured relative to the credible reaction of dropping unprofitable policies. For the single-policy situation, this assumption turns the best pooling policy into an equilibrium when it dominates the Walrasian policy, and thus obviates the nonexistence problem. However, the pooling policy is never an equilibrium when firms are allowed to offer more than one policy.

Finally, both the Walrasian and the pooling policies are Pareto-dominated by the subsidy equilibrium, when the latter differs from the Walrasian policy.

In the many-type case, we focus on a particular subclass of the possible reaction strategies. It is credible and has the additional advantage of being easy to interpret: Firms that lose money will drop all their policies. There is a unique Nash equilibrium portfolio supported by these reactions, and it has a convenient hierarchical optimality property.

Within the context of this model, several open questions remain. First, it is possible that there may be many equilibria relative to general reaction strategies of the sort defined in Sec. V. Are they different, and if so, how do these differences show up in the nature of the optimality problem solved by the equilibria? This question is interesting because it throws into sharp relief the disaggregated aspects of the interplay between efficiency and equilibrium.

Second, it would be interesting to investigate whether the results change if players evaluate strategies according to the present discounted value of a stream of profits rather than the profit earned after all adjustments have been completed. This is an essential next

step if the model is to be broadened to include the consumers as strategic players, and to allow firms to use such reasonable learning tools as experience-rating.

A third and related question has to do with the industrial organization of the insurance market. In carrying out the many-type analysis, we assumed that there was a single active firm. While we partially justified this assumption in the context of the reactions we employed, it seems clear that a more general approach would allow for firms to offer subsets of a full portfolio of policies, as long as each subset broke even. Under these circumstances, the credibility of a reaction like the one we use may be threatened, and other reactions might be needed. It might turn out to be the case that the perfect equilibrium conditions tell us something about the industrial structure.

REFERENCES

Cave, J.A.K., "Perfect Equilibrium in Continuous Games," January 1984 (mimeograph).

Dasgupta, P., and E. Maskin, "The Existence of Equilibrium in Discontinuous Games": Part 2, "Applications," June 1982 (mimeograph).

Hoy, M., "Categorizing Risks in the Insurance Industry," *Quarterly Journal of Economics,* Vol. 46, May 1982, pp. 321-336.

Miyazaki, H., "The Rat Race and Internal Labor Markets," *Journal of Economics,* Vol. 8, August 1977, pp. 394-418.

Marshak, T., and R. Selten, "Restabilizing Responses, Inertia Supergames, and Oligopolistic Equilibria," *Quarterly Journal of Economics,* February 1978, pp. 71-93.

Riley, J., "Informational Equilibria," *Econometrica,* Vol. 47, March 1979, pp. 331-360.

Rosenthal, R., and A. Weiss, "Mixed-Strategy Equilibria in Markets with Asymmetric Information," 1983 (mimeograph).

Rothschild, M., and J. Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics,* Vol. 40, November 1976, pp. 629-649.

Selten, R., "A Reexamination of the Perfectness Concept for Games in Extensive Form," *International Journal of Game Theory,* Vol. 2, 1973, pp. 25-55.

Spence, M., "Product Differentiation and Performance in Insurance Markets," *Journal of Public Economics,* Vol. 10, 1978, pp. 427-447.

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

RAND/R-3015-HHS