

EQUILIBRIUM IN COMPETITIVE INSURANCE MARKETS: AN ESSAY ON THE ECONOMICS OF IMPERFECT INFORMATION*

MICHAEL ROTHSCHILD AND JOSEPH STIGLITZ

Introduction, 629.—I. The basic model, 630.—II. Robustness, 638—III. Conclusion, 648.

INTRODUCTION

Economic theorists traditionally banish discussions of information to footnotes. Serious consideration of costs of communication, imperfect knowledge, and the like would, it is believed, complicate without informing. This paper, which analyzes competitive markets in which the characteristics of the commodities exchanged are not fully known to at least one of the parties to the transaction, suggests that this comforting myth is false. Some of the most important conclusions of economic theory are not robust to considerations of imperfect information.

We are able to show that not only may a competitive equilibrium not exist, but when equilibria do exist, they may have strange properties. In the insurance market, upon which we focus much of our discussion, sales offers, at least those that survive the competitive process, do not specify a price at which customers can buy all the insurance they want, but instead consist of both a price and a quantity—a particular amount of insurance that the individual can buy at that price. Furthermore, if individuals were willing or able to reveal their information, everybody could be made better off. By their very being, high-risk individuals cause an externality: the low-risk individuals are worse off than they would be in the absence of the high-risk individuals. However, the high-risk individuals are no better off than they would be in the absence of the low-risk individuals.

These points are made in the next section by analysis of a simple model of a competitive insurance market. We believe that the lessons gleaned from our highly stylized model are of general interest, and attempt to establish this by showing in Section II that our model is robust and by hinting (space constraints prevent more) in the conclusion that our analysis applies to many other situations.

* This work was supported by National Science Foundation Grants SOC 74-22182 at the Institute for Mathematical Studies in the Social Sciences, Stanford University and SOC 73-05510 at Princeton University. The authors are indebted to Steve Salop, Frank Hahn, and Charles Wilson for helpful comments, and to the participants in the seminars at the several universities at which these ideas were presented.

I. THE BASIC MODEL

Most of our argument can be made by analysis of a very simple example. Consider an individual who will have an income of size W if he is lucky enough to avoid accident. In the event an accident occurs, his income will be only $W - d$. The individual can insure himself against this accident by paying to an insurance company a premium α_1 , in return for which he will be paid $\hat{\alpha}_2$ if an accident occurs. Without insurance his income in the two states, "accident," "no accident," was $(W, W - d)$; with insurance it is now $(W - \alpha_1, W - d + \alpha_2)$, where $\alpha_2 = \hat{\alpha}_2 - \alpha_1$. The vector $\alpha = (\alpha_1, \alpha_2)$ completely describes the insurance contract.¹

1.1 Demand for Insurance Contracts

On an insurance market, insurance contracts (the α 's) are traded. To describe how the market works, it is necessary to describe the supply and demand functions of the participants in the market. There are only two kinds of participants, individuals who buy insurance and companies that sell it. Determining individual demand for insurance contracts is straightforward. An individual purchases an insurance contract so as to alter his pattern of income across states of nature. Let W_1 denote his income if there is no accident and W_2 his income if an accident occurs; the expected utility theorem states that under relatively mild assumptions his preferences for income in these two states of nature are described by a function of the form,

$$(1) \quad \hat{V}(p, W_1, W_2) = (1 - p)U(W_1) + pU(W_2),$$

where $U(\)$ represents the utility of money income² and p the probability of an accident. Individual demands may be derived from (1). A contract α is worth $V(p, \alpha) = \hat{V}(p, W - \alpha_1, W - d + \alpha_2)$. From

1. Actual insurance contracts are more complicated because a single contract will offer coverage against many potential losses. A formal generalization of the scheme above to cover this case is straightforward. Suppose that an individual will, in the absence of insurance, have an income of W_i if state i occurs. An insurance contract is simply an n -tuple $(\alpha_1, \dots, \alpha_n)$ whose i -th coordinate describes the net payment of the individual to the insurance company if state i occurs. We confine our discussion to the simple case mentioned in the text, although it could be trivially extended to this more complicated case.

Many insurance contracts are not as complicated as the n -tuples described above—Blue Cross schedules listing maximum payments for specific illnesses and operations are an isolated example—but are instead resolvable into a fixed premium and a payment schedule that is in general a simple function of the size of the loss such as $F(L) = \text{Max} [0, c(L-D)]$, where $c \times 100\%$ is the co-insurance rate and D is the deductible. With such a contract when a loss occurs, determining its size is often a serious problem. In other words, finding out exactly what state of the world has occurred is not always easy. We ignore these problems. A large literature analyzes optimal insurance contracts. See, for example, Arrow (1971) and Borch (1968).

2. We assume that preferences are not state-dependent.

all the contracts the individual is offered, he chooses the one that maximizes $V(p, \alpha)$. Since he always has the option of buying no insurance, an individual will purchase a contract α only if $V(p, \alpha) \geq V(p, 0) = \hat{V}(p, W, W - d)$. We assume that persons are identical in all respects save their probability of having an accident and that they are risk-averse ($U'' < 0$); thus $V(p, \alpha)$ is quasi-concave.

1.2 Supply of Insurance Contracts

It is less straightforward to describe how insurance companies decide which contracts they should offer for sale and to which people. The return from an insurance contract is a random variable. We assume that companies are risk-neutral, that they are concerned only with expected profits, so that contract α when sold to an individual who has a probability of incurring an accident of p , is worth

$$(2) \quad \pi(p, \alpha) = (1 - p)\alpha_1 - p\alpha_2 = \alpha_1 - p(\alpha_1 + \alpha_2).$$

Even if firms are not expected profit maximizers, on a well-organized competitive market they are likely to behave as if they maximized (2).³

Insurance companies have financial resources such that they are willing and able to sell any number of contracts that they think will make an expected profit.⁴ The market is competitive in that there is free entry. Together these assumptions guarantee that any contract that is demanded and that is expected to be profitable will be supplied.

3. Since the theory of the firm behavior under uncertainty is one of the more unsettled areas of economic theory, we cannot look to it for the sort of support of any assumption we might make, which the large body of literature devoted to the expected utility theorem provides for equation (1) above. Nonetheless, two arguments (and the absence of a remotely as attractive distinguishable alternative) justify (2): the first is the rather vaguely supported but widely held proposition that companies owned by stockholders who themselves hold diversified portfolios ought to maximize their expected profits; management that does not follow this policy will be displaced. The second supposes that insurance companies are held by a large number of small shareholders each of whom receives a small share of the firm's profits. If the risks insured against are independent or otherwise diversifiable, then the law of large numbers guarantees that each shareholder's return will be approximately constant and any individual insurance contract contributes to his profits only through its expected value. In this case stockholders' interests will be well served if, and only if, management maximizes expected profits.

A variant of the second argument is obtained by considering the case in which shareholders and policyholders are the same people, or in more familiar terms, when the insurance company is a mutual company. In this case the insurance company is just a mechanism for risk pooling. Under conditions where diversification is possible, each contract's contribution to the company's dividend (or loss) is proportional to its expected value.

4. The same kinds of arguments used to justify (2)—in particular the appeal to the law of large numbers—can be used to justify this assumption. Weaker conditions than independence will suffice. See Revesz (1960), p. 190, for a theorem that states roughly that, if insurance contracts can be arranged in space so that even though con-

1.3 Information about Accident Probabilities

We have not so far discussed how customers and companies come to know or estimate the parameter p , which plays such a crucial role in the valuation formulae (1) and (2). We make the bald assumption that individuals know their accident probabilities, while companies do not. Since insurance purchasers are identical in all respects save their propensity to have accidents, the force of this assumption is that companies cannot discriminate among their potential customers on the basis of their characteristics. This assumption is defended and modified in subsection II.1.

A firm may use its customers' market behavior to make inferences about their accident probabilities. Other things equal, those with high accident probabilities will demand more insurance than those who are less accident-prone. Although possibly accurate, this is not a profitable way of finding out about customer characteristics. Insurance companies want to know their customers' characteristics in order to decide on what terms they should offer to let them buy insurance. Information that accrues after purchase may be used only to lock the barn after the horse has been stolen.

It is often possible to force customers to make market choices in such a way that they both reveal their characteristics and make the choices the firm would have wanted them to make had their characteristics been publicly known. In their contribution to this symposium, Salop and Salop call a market device with these characteristics a *self-selection mechanism*. Analysis of the functioning of self-selection mechanisms on competitive markets is a major focus of this paper.

1.4 Definition of Equilibrium

We assume that customers can buy only one insurance contract. This is an objectionable assumption. It implies, in effect, that the seller of insurance specifies both the prices and quantities of insurance purchased. In most competitive markets, sellers determine only price and have no control over the amount their customers buy. Nonetheless, we believe that what we call price and quantity competition is more appropriate for our model of the insurance market than tradi-

tracts that are close to one another are not independent, those that are far apart are approximately independent, then the average return from all contracts is equal to its expected value with probability one. Thus, an insurance company that holds a large number of health policies should be risk-neutral, even though the fact that propinquity carries illness implies that not all insured risks are independent. Some risks that cannot be diversified; i.e., the risk of nuclear war (or of a flood or a plague) cannot be spread by appeal to the law of large numbers. Our model applies to diversifiable risks. This class of risks is considerably larger than the independent ones.

tional price competition. We defend this proposition at length in subsection II.2 below.

Equilibrium in a competitive insurance market is a set of contracts such that, when customers choose contracts to maximize expected utility, (i) no contract in the equilibrium set makes negative expected profits; and (ii) there is no contract outside the equilibrium set that, if offered, will make a nonnegative profit. This notion of equilibrium is of the Cournot-Nash type; each firm assumes that the contracts its competitors offer are independent of its own actions.

1.5 Equilibrium with Identical Customers

Only when customers have different accident probabilities, will insurance companies have imperfect information. We examine this case below. To illustrate our, mainly graphical, procedure, we first analyze the equilibrium of a competitive insurance market with identical customers.⁵




FIGURE I

In Figure I the horizontal and vertical axes represent income in

5. The analysis is identical if individuals have different p 's, but companies know the accident probabilities of their customers. The market splits into several submarkets—one for each different p represented. Each submarket has the equilibrium described here.

the states: no accident, accident, respectively. The point E with coordinates (\hat{W}_1, \hat{W}_2) is the typical customer's uninsured state. Indifference curves are level sets of the function of equation (1). Purchasing the insurance policy $\alpha = (\alpha_1, \alpha_2)$ moves the individual from E to the point $(\hat{W}_1 - \alpha_1, \hat{W}_2 + \alpha_2)$.

Free entry and perfect competition will ensure that policies bought in competitive equilibrium make zero expected profits, so that if α is purchased,

$$(3) \quad \alpha_1(1 - p) - \alpha_2 p = 0.$$

The set of all policies that break even is given analytically by (3) and diagrammatically by the line EF in Figure I, which is sometimes referred to as the fair-odds line. The equilibrium policy α^* maximizes the individual's (expected) utility and just breaks even. Purchasing α^* locates the customer at the tangency of the indifference curve with the fair-odds line. α^* satisfies the two conditions of equilibrium: (i) it breaks even; (ii) selling any contract preferred to it will bring insurance companies expected losses.

Since customers are risk-averse, the point α^* is located at the intersection of the 45°-line (representing equal income in both states of nature) and the fair-odds line. In equilibrium each customer buys complete insurance at actuarial odds. To see this, observe that the slope of the fair-odds line is equal to the ratio of the probability of not having an accident to the probability of having an accident $((1 - p)/p)$, while the slope of the indifference curve (the marginal rate of substitution between income in the state no accident to income in the state accident) is $[U'(W_1)(1 - p)]/[U'(W_2)p]$, which, when income in the two states is equal, is $(1 - p)/p$, independent of U .

1.6 Imperfect Information: Equilibrium with Two Classes of Customers

Suppose that the market consists of two kinds of customers: low-risk individuals with accident probability p^L , and high-risk individuals with accident probability $p^H > p^L$. The fraction of high-risk customers is λ , so the average accident probability is $\bar{p} = \lambda p^H + (1 - \lambda)p^L$. This market can have only two kinds of equilibria: *pooling equilibria* in which both groups buy the same contract, and *separating equilibria* in which different types purchase different contracts.

A simple argument establishes that *there cannot be a pooling equilibrium*. The point E in Figure II is again the initial endowment of all customers. Suppose that α is a pooling equilibrium and consider $\pi(\bar{p}, \alpha)$. If $\pi(\bar{p}, \alpha) < 0$, then firms offering α lose money, contradicting




FIGURE II

the definition of equilibrium. If $\pi(\bar{p}, \alpha) > 0$, then there is a contract that offers slightly more consumption in each state of nature, which still will make a profit when all individuals buy it. All will prefer this contract to α , so α cannot be an equilibrium. Thus, $\pi(\bar{p}, \alpha) = 0$, and α lies on the market odds line EF (with slope $(1 - \bar{p})/\bar{p}$).

It follows from (1) that at α the slope of the high-risk indifference curve through α , \bar{U}^H , is $(p^L/1 - p^L)(1 - p^H/p^H)$ times the slope of \bar{U}^L , the low-risk indifference curve through α . In this figure \bar{U}^H is a broken line, and \bar{U}^L a solid line. The curves intersect at α ; thus there is a contract, β in Figure II, near α , which low-risk types prefer to α . The high risk prefer α to β . Since β is near α , it makes a profit when the less risky buy it, $(\pi(p^L, \beta) \simeq \pi(p^L, \alpha) > \pi(\bar{p}, \alpha) = 0)$. The existence of β contradicts the second part of the definition of equilibrium; α cannot be an equilibrium.

If there is an equilibrium, each type must purchase a separate contract. Arguments, which are, we hope, by now familiar, demonstrate that each contract in the equilibrium set makes zero profits. In Figure III the low-risk contract lies on line EL (with slope $(1 - p^L)/p^L$), and the high-risk contract on line EH (with slope $(1 - p^H)/p^H$). As was shown in the previous subsection, the contract on EH most preferred by high-risk customers gives complete insurance.




FIGURE III

This is α^H in Figure III; it must be part of any equilibrium. Low-risk customers would, of all contracts on EL , most prefer contract β which, like α^H , provides complete insurance. However, β offers more consumption in each state than α^H , and high-risk types will prefer it to α^H . If β and α^H are marketed, both high- and low-risk types will purchase β . The nature of imperfect information in this model is that insurance companies are unable to distinguish among their customers. All who demand β must be sold β . Profits will be negative; (α^H, β) is not an equilibrium set of contracts.

An equilibrium contract for low-risk types must not be more attractive to high-risk types than α^H ; it must lie on the southeast side of U^H , the high-risk indifference curve through α^H . We leave it to the reader to demonstrate that of all such contracts, the one that low-risk types most prefer is α^L , the contract at the intersection of EL and U^H in Figure III. This establishes that *the set (α^H, α^L) is the only possible equilibrium for a market with low- and high-risk customers.*⁶ However, (α^H, α^L) may not be an equilibrium. Consider the contract γ in Figure III. It lies above U^L , the low-risk indifference curve through α^L and also above U^H . If γ is offered, both low- and high-risk types

6. This largely heuristic argument can be made completely rigorous. See Wilson (1976).

will purchase it in preference to either α^H or α^L . If it makes a profit when both groups buy it, γ will upset the potential equilibrium of (α^H, α^L) . γ 's profitability depends on the composition of the market. If there are sufficiently many high-risk people that EF represents market odds, then γ will lose money. If market odds are given by EF' (as they will be if there are relatively few high-risk insurance customers), then γ will make a profit. Since (α^H, α^L) is the only possible equilibrium, in this case the competitive insurance market will have no equilibrium.

This establishes that *a competitive insurance market may have no equilibrium.*

We have not found a simple intuitive explanation for this non-existence; but the following observations, prompted by Frank Hahn's note (1974), may be suggestive. The information that is revealed by an individual's choice of an insurance contract depends on all the other insurance policies offered; there is thus a fundamental informational externality that each company, when deciding on which contract it will offer, fails to take into account. Given any set of contracts that breaks even, a firm may enter the market using the informational structure implicit in the availability of that set of contracts to make a profit; at the same time it forces the original contracts to make a loss. But as in any Nash equilibrium, the firm fails to take account of the consequences of its actions, and in particular, the fact that when those policies are no longer offered, the informational structure will have changed and it can no longer make a profit.

We can characterize the conditions under which an equilibrium does not exist. An equilibrium will not exist if the costs to the low-risk individual of pooling are low (because there are relatively few of the high-risk individuals who have to be subsidized, or because the subsidy per individual is low, i.e., when the probabilities of the two groups are not too different), or if their costs of separating are high. The costs of separating arise from the individual's inability to obtain complete insurance. Thus, the costs of separating are related to the individuals' attitudes toward risk. Certain polar cases make these propositions clear. If $p^L = 0$, it never pays the low-risk individuals to pool, and by continuity, for sufficiently small p^L it does not pay to pool. Similarly, if individuals are risk-neutral, it never pays to pool; if they are infinitely risk averse with utility functions

$$(1') \quad \bar{V}(p, W_1, W_2) = \text{Min}(W_1, W_2),$$

it always pays to pool.

1.7 Welfare Economics of Equilibrium

One of the interesting properties of the equilibrium is that the presence of the high-risk individuals exerts a negative externality on the low-risk individuals. The externality is completely dissipative; there are losses to the low-risk individuals, but the high-risk individuals are no better off than they would be in isolation.

If only the high-risk individuals would admit to their having high accident probabilities, all individuals would be made better off without anyone being worse off.

The separating equilibrium we have described may not be Pareto optimal even relative to the information that is available. As we show in subsection II.3 below, there may exist a pair of policies that break even together and that make both groups better off.

II. ROBUSTNESS

The analysis of Section I had three principal conclusions: First, competition on markets with imperfect information is more complex than in standard models. Perfect competitors may limit the quantities their customers can buy, not from any desire to exploit monopoly power, but simply in order to improve their information. Second, equilibrium may not exist. Finally, competitive equilibria are not Pareto optimal. It is natural to ask whether these conclusions (particularly the first, which was an assumption rather than a result of the analysis) can be laid to the special and possibly strained assumptions of our model. We think not. Our conclusions (or ones very like) must follow from a serious attempt to comprehend the workings of competition with imperfect and asymmetric information. We have analyzed the effect of changing our model in many ways. The results were always essentially the same.

Our attempts to establish robustness took two tacks. First, we showed that our results did not depend on the simple technical specifications of the model. This was tedious, and we have excised most of the details from the present version. The reader interested in analysis of the effects (distinctly minor) of changing our assumptions that individuals are alike in all respects save their accident probabilities, that there are only two kinds of customers, and that the insurance market lasts but a single period, is referred to earlier versions of this paper.⁷ An assessment of the importance of the as-

7. See Rothschild and Stiglitz (1975). One curious result of these investigations should be mentioned. In other areas of economic theory where existence of equilibrium has been a problem, smoothing things by introducing a continuum of individuals of

sumption that individuals know their accident probabilities, while insurance companies do not (which raises more interesting issues), is given in subsection II.1 below.

Another approach to the question of robustness is the subject of the next three subsections. In them we question the behavioral assumptions and the equilibrium concepts used in Section I.

II.1 Information Assumptions

Suppose that there are two groups of customers and that not all individuals within each group have the same accident probability. The average accident probability of one group is greater than that of the other; individuals within each group know the mean accident probability for members of their group, but do not know their own accident probabilities. As before, the insurance company cannot tell directly the accident probability of any particular individual, or even the group to which he belongs. For example, suppose that some persons occasionally drink too much, while the others almost never drink. Insurance firms cannot discover who drinks and who does not. Individuals know that drinking affects accident probabilities, but it affects different people differently. Each individual does not know how it will affect him.

In such a situation the expected utility theorem states that individuals make (and behave according to) estimates of their accident probabilities; if these estimates are unbiased in the sense that the average accident probability of those who estimate their accident probability to be p actually is p , then the analysis goes through as before.

Unbiasedness seems a reasonable assumption (what is a more attractive alternative?). However, not even this low level of correctness of beliefs is required for our conclusions. Suppose, for example, that individuals differ both with respect to their accident probabilities and to their risk aversion, but they all assume that their own accident probabilities are \bar{p} . If low-risk individuals are less risk-averse on average, then there will not exist a pooling equilibrium; there may exist no equilibrium at all; and if there does exist an equilibrium, it will entail partial insurance for both groups. Figure IV shows that there

different types can insure existence. Not so here. If there is a continuous distribution of accident probabilities (but customers are otherwise identical), then equilibrium never exists. There is an intuitive explanation for this striking result. We argued above that, if accident probabilities were close together, then equilibrium would not exist. When there is a continuum of probabilities, there always are individuals with close probabilities with whom it pays to "pool." For a proof of this result, which is not elementary, see Riley (1976).




FIGURE IV

will not exist a pooling equilibrium. If there were a pooling equilibrium, it would clearly be with complete insurance at the market odds, since both groups' indifference curves have the slope of the market odds line there. If the low-risk individuals are less risk-averse, then the two indifference curves are tangent at F , but elsewhere the high-risk individuals' indifference curve lies above the low-risk individuals' indifference curve. Thus, any policy in the shaded area between the two curves will be purchased by the low-risk individuals in preference to the pooling contract at F .

Other such cases can be analyzed, but we trust that the general principle is clear. Our pathological conclusions do not require that people have particularly good information about their accident probabilities. They will occur under a wide variety of circumstances, including the appealing case of unbiasedness. Neither insurance firms nor their customers have to be perfectly informed about the differences in risk properties that exist among individuals: What is required is that individuals with different risk properties differ in some characteristic that can be linked with the purchase of insurance and that, somehow, insurance firms discover this link.

II.2 Price Competition Versus Quantity Competition

One can imagine our model of the insurance market operating in two distinct modes. The first, price competition, is familiar to all

students of competitive markets. Associated with any insurance contract α is a number $q(\alpha) = \alpha_1/\alpha_2$, which, since it is the cost per unit coverage, is called the price of insurance. Under price competition, insurance firms establish a price of insurance and allow their customers to buy as much or as little insurance as they want at that price. Thus, if contract α is available from a company, so are the contracts 2α and $(1/2)\alpha$; the former pays twice as much benefits (and costs twice as much in premiums) as α ; the latter is half as expensive and provides half as much coverage.

Opposed to price competition is what we call price and quantity competition. In this regime companies may offer a number of different contracts, say $\alpha^1, \alpha^2, \dots, \alpha^n$. Individuals may buy at most one contract. They are not allowed to buy arbitrary multiples of contracts offered, but must instead settle for one of the contracts explicitly put up for sale. A particular contract specifies both a price and a quantity of insurance. Under price and quantity competition it is conceivable that insurance contracts with different prices of insurance will exist in equilibrium; people who want more insurance may be willing to pay a higher price for it (accept less favorable odds) than those who make do with shallower coverage. Under price competition customers will buy insurance only at the lowest price quoted in the market.

The argument of Section I depends heavily on our assumption that price and quantity competition, and not simply price competition, characterizes the competitive insurance market. This assumption is defended here. The argument is basically quite simple. Price competition is a special case of price and quantity competition. Nothing in the definition of price and quantity competition prevents firms from offering for sale a set of contracts with the same price of insurance. Since the argument above characterized all equilibria under price and quantity competition, it also characterized all equilibria when some firms set prices and others set prices and quantities. Thus, it must be that price competition cannot compete with price and quantity competition.⁸

This argument hinges on one crucial assumption: regardless of the form of competition, customers purchase but a single insurance contract or equivalently that the total amount of insurance purchased

8. We leave to the reader a detailed proof. A sketch follows. Suppose that there are two groups in the population. If the price of insurance is q , high- and low-risk customers will buy $\alpha^H(q)$ and $\alpha^L(q)$, respectively. It is easy to figure out what total insurance company profits, $P(q)$, are. The equilibrium price q^* is the smallest q such that $P(q) = 0$. Since $P(q)$ is continuous in q and it is easy to find q such that $P(q) > 0$ and $P(q) < 0$, such a q^* exists. To show that price competition will not survive, it is only necessary to show that $(\alpha^H(q^*), \alpha^L(q^*))$ is not an equilibrium set of contracts as defined in subsection I.4 above.

by any one customer is known to all companies that sell to him. We think that this is an accurate description of procedures on at least some insurance markets. Many insurance policies specify either that they are not in force if there is another policy or that they insure against only the first, say, \$1,000 of losses suffered. That is, instead of being a simple bet for or against the occurrence of a particular event, an insurance policy is a commitment on the part of the company to restore at least partially the losses brought about by the occurrence of that event. The person who buys two \$1,000 accident insurance policies does not have \$2,000 worth of protection. If an accident occurs, all he gets from his second policy is the privilege of watching his companies squabble over the division of the \$1,000 payment. There is no point in buying more than one policy.

Why should insurance markets operate in this way? One simple and obvious explanation is moral hazard. Because the insured can often bring about, or at least make more likely, the event being insured against, insurance companies want to limit the amount of insurance their customers buy. Companies want to see that their customers do not purchase so much insurance that they have an interest in an accident occurring. Thus, companies will want to monitor the purchases of their customers. Issuing contracts of the sort described above is the obvious way to do so.

A subtler explanation for this practice is provided by our argument that price and quantity competition can dominate price competition. If the market is in equilibrium under price competition, a firm can offer a contract, specifying price and quantity, that will attract the low-risk customers away from the companies offering contracts specifying price alone. Left with only high-risk customers, these firms will lose money. This competitive gambit will successfully upset the price competition equilibria if the entering firm can be assured that those who buy its contracts hold no other insurance. Offering insurance that pays off only for losses not otherwise insured is a way to guarantee this.

It is sometimes suggested that the term "competitive" can be applied only to markets where there is a single price of a commodity and each firm is a price taker. This seems an unnecessarily restrictive use of the term competitive. The basic idea underlying competitive markets involves free entry and noncollusive behavior among the participants in the market. In some economic environments price taking without quantity restrictions is a natural result of such markets. In the situations described in this paper, this is not so.

II.3 Restrictions on Firm Behavior and Optimal Subsidies

An important simplification of the analysis of Section I was the assumption that each insurance company issued but a single contract. We once thought this constraint would not affect the nature of equilibrium. We argued that in equilibrium firms must make nonnegative profits. Suppose that a firm offers two contracts, one of which makes an expected profit of say, $\$S$, per contract sold, the other an expected loss of $\$L$ per contract. The firm can make nonnegative expected profits if the ratio of the profitable to the unprofitable contracts sold is at least μ , where $\mu = L/S$. However, the firm can clearly make more profits if it sells only the contracts on which it makes a profit. It and its competitors have no reason to offer the losing contracts, and in competitive equilibrium, they will not be offered. Since only contracts that make nonnegative profits will be offered, it does not matter, given our assumptions about entry, that firms are assumed to issue only a single contract. If there is a contract that could make a profit, a firm will offer it.

This argument is not correct. The possibility of offering more than one contract is important to firms, and to the nature and existence of equilibrium. Firms that offer several contracts are not dependent on the policies offered by other firms for the information generated by the choices of individuals. By offering a menu of policies, insurance firms may be able to obtain information about the accident probabilities of particular individuals. Furthermore, although there may not be an equilibrium in which the profits from one contract subsidize the losses of another contract, it does not follow that such a pair of contracts cannot break what would otherwise be an equilibrium.

Such a case is illustrated in Figure V. EF is again the market odds line. A separating equilibrium exists $(\bar{\alpha}^H, \bar{\alpha}^L)$. Suppose that a firm offered the two contracts, $\alpha^{H'}$ and $\alpha^{L'}$; $\alpha^{H'}$ makes a loss, $\alpha^{L'}$ makes a profit. High-risk types prefer $\alpha^{H'}$ to $\bar{\alpha}^H$, and low-risk types prefer $\alpha^{L'}$ to $\bar{\alpha}^L$. These two contracts, if offered by a single firm together, do not make losses. The profits from $\alpha^{L'}$ subsidize the losses of $\alpha^{H'}$. Thus, $(\alpha^{H'}, \alpha^{L'})$ upsets the equilibrium $(\bar{\alpha}^H, \bar{\alpha}^L)$.

This example points up another possible inefficiency of separating equilibria. Consider the problem of choosing two contracts (α^H, α^L) such that α^L maximizes the utility of the low-risk individual subject to the constraints that (a) the high-risk individual prefers α^H to α^L and (b) the pair of contracts α^H and α^L break even when bought by high- and low-risk types, respectively, in the ratio λ to $(1 - \lambda)$. This




FIGURE V

is a kind of optimal subsidy problem. If the separating equilibrium, when it exists, does not solve this problem, it is inefficient. Figure V shows that the separating equilibrium can be inefficient in this sense. We now show that if there are enough high-risk people, then the separating equilibrium can be efficient.

The optimal subsidy problem always has a solution $(\alpha^{H*}, \alpha^{L*})$. The optimal high-risk contract α^{H*} will always entail complete insurance so that $V(p^H, \alpha^{H*}) = U(W - p^H d + a)$, where a is the per capita subsidy of the high risk by the low risk. This subsidy decreases income for each low-risk person by γa (where $\gamma = \lambda/(1 - \lambda)$) in each state. Net of this charge α^{L*} breaks even when low-risk individuals buy it. Thus, $\alpha^{L*} = (\alpha_1 + \gamma a, \alpha_2 - \gamma a)$, where $\alpha_1 = \alpha_2 p^L / (1 - p^L)$.

To find the optimal contract, one solves the following problem: Choose a and α_2 to maximize

$$U(X) (1 - p^L) + U(Z)p^L,$$

subject to

$$U(Y) \geq U(X) (1 - p^H) + U(Z)p^H$$

$$a \geq 0,$$

where

$$X = W_0 - \gamma a - \alpha_2 p^L / (1 - p^L),$$

$$Y = W_0 - p^H d + a,$$

and

$$Z = W_0 - d - \gamma a + \alpha_2.$$

The solution to this problem can be analyzed by standard Kuhn-Tucker techniques. If the constraint $a \geq 0$ is binding at the optimum, then the solution involves no subsidy to the high-risk persons; $(\alpha^{H*}, \alpha^{L*})$ is the separating equilibrium. It is straightforward but tedious to show that a sufficient condition for this is that

$$(4) \quad \frac{(p^H - p^L) \gamma}{p^L(1 - p^L)} > \frac{U'(Y)[U'(Z) - U'(X)]}{U'(X)U'(Z)}$$

where X , Y , and Z are determined by the optimal a^* , α_2^* . The right-hand side of (4) is always less than

$$\frac{U'(W_0 - d)[U'(W_0 - d) - U'(W_0)]}{U'(W_0)^2}$$

so that there exist values of γ (and thus of λ) large enough to satisfy (4).

II.4 Alternative Equilibrium Concepts

There are a number of other concepts of equilibrium that we might have employed. These concepts differ with respect to assumptions concerning the behavior of the firms in the market. In our model the firm assumes that its actions do not affect the market—the set of policies offered by other firms was independent of its own offering.

In this subsection we consider several other equilibrium concepts, implying either less or more rationality in the market. We could, for instance, call any set of policies that just break even given the set of individuals who purchase them an *informationally consistent equilibrium*. This assumes that the forces for the creation of new contracts are relatively weak (in the absence of profits). Thus, in Figure III, α^H and any contract along the line EL below α^L is a set of informationally consistent separating equilibrium contracts; any single contract along the line EF is an informationally consistent pooling equilibrium contract. This is the notion of equilibrium that Spence (1973) has employed in most of his work. The longer the lags in the system, the greater the difficulty of competing by offering different contracts, the more stable is an informationally consistent equilibrium. Thus, while

this seems to us a reasonable equilibrium concept for the models of educational signaling on which Spence focused, it is less compelling when applied to insurance or credit markets (see Jaffee and Russell's contribution to this symposium).

A *local* equilibrium is a set of contracts such that there do not exist any contracts in the vicinity of the equilibrium contracts that will be chosen and make a positive profit. If we rule out the subsidies of the last subsection, then the set of separating contracts, which maximizes the welfare of low-risk individuals, is a local equilibrium.

The notion that firms experiment with contracts similar to those already on the market motivates the idea of a local equilibrium. Even if firms have little knowledge about the shape of utility functions, and about the proportions of population in different accident probabilities, one would expect that competition would lead to small perturbations around the equilibrium. A stable equilibrium requires that such perturbations not lead to firms making large profits, as would be the case with some perturbations around a pooling point.

These two concepts of equilibrium imply that firms act less rationally than we assumed they did in Section I. It is possible that firms exhibit a greater degree of rationality; that is, firms ought not to take the set of contracts offered by other firms as given, but ought to assume that other firms will act as they do, or at least will respond in some way to the new contract offered by the firm. Hence, in those cases where in our definition there was no equilibrium, because for any set of contracts there is a contract that will break even and be chosen by a subset of the population, given that the contracts offered by the other firms remain unchanged, those contracts that break the equilibrium may not break even if the other firms also change their contracts. The peculiar provision of many insurance contracts, that the effective premium is not determined until the *end* of the period (when the individual obtains what is called a dividend), is perhaps a reflection of the uncertainty associated with who will purchase the policy, which in turn is associated with the uncertainty about what contracts other insurance firms will offer.

Wilson (1976) introduced and analyzed one such nonmyopic equilibrium concept. A Wilson equilibrium is a set of contracts such that, when customers choose among them so as to maximize profits, (a) all contracts make nonnegative profits and (b) there does not exist a new contract (or set of contracts), which, if offered, makes positive profits even when all contracts that lose money as a result of this entry are withdrawn. In the simple model of Section I, such equilibria always

exist. Comparing this definition with the one of subsection I.4 above makes it clear that, when it exists, our separating equilibrium is also a Wilson equilibrium. When this does not exist, the Wilson equilibrium is the pooling contract that maximizes the utility of the low-risk customers. This is β in Figure VI. β dominates the separating pair (α^L, α^H) . Consider a contract like γ , which the low risk prefer to β . Under our definition of equilibrium it upsets β . Under Wilson's it does not. When the low risk desert β for γ , it loses money and is withdrawn. Then the high risk also buy γ . When both groups buy γ , it loses money. Thus, γ does not successfully compete against β .

Although this equilibrium concept is appealing, it is not without its difficulties. It seems a peculiar halfway house; firms respond to competitive entry by dropping policies, but not by adding new policies. Furthermore, although counterexamples are very complicated to construct, it appears that a Wilson equilibrium may not exist if groups differ in their attitudes towards risk. Finally, in the absence of collusion or regulation, in a competitive insurance market, it is hard to see how or why any single firm should take into account the consequences of its offering a new policy. On balance, it seems to us that nonmyopic equilibrium concepts are more appropriate for models of monopoly (or oligopoly) than for models of competition.




FIGURE VI

III. CONCLUSION

We began this research with the hope of showing that even a small amount of imperfect information could have a significant effect on competitive markets. Our results were more striking than we had hoped: the single price equilibrium of conventional competitive analysis was shown to be no longer viable; market equilibrium, when it existed, consisted of contracts which specified both prices and quantities; the high-risk (low ability, etc.) individuals exerted a dissipative externality on the low-risk (high ability) individuals; the structure of the equilibrium as well as its existence depended on a number of assumptions that, with perfect information, were inconsequential; and finally, and in some ways most disturbing, under quite plausible conditions equilibrium did not exist.

Our analysis, and our conclusions, extend beyond the simple insurance market described above. The models of educational screening and signaling studied by, among others, Arrow (1973), Riley (1975), Spence (1973, 1974), and Stiglitz (1971, 1972, 1974a, 1975b), are obvious examples. The other papers in this symposium describe models that can be profitably studied using our techniques and our concepts. Models in which communities choose the level of public goods and individuals choose among communities on the basis of the menu of public goods and taxes that the different communities offer, provide a less obvious but, we think, important case.⁹

Do these theoretical speculations tell us anything about the real world? In the absence of empirical work it is hard to say. The market on which we focused most of our analysis, that for insurance, is probably not competitive; whether our model may partially explain this fact is almost impossible to say. But there are other markets, particularly financial and labor markets, which appear to be competitive and in which imperfect and asymmetric information play an important role. We suspect that many of the peculiar institutions of these labor markets arise as responses to the difficulties that they, or any competitive market, have in handling problems of information. Establishing (or refuting) this conjecture seems to provide a rich agenda for future research.

UNIVERSITY OF WISCONSIN, MADISON

STANFORD UNIVERSITY AND ALL SOULS COLLEGE, OXFORD

9. See F. Westhoff's dissertation (1974), and Stiglitz (1974b). A more complete discussion of these is in our earlier working paper referred to in footnote 7 above. Salop and Salop (1972) demonstrated, in an early draft of their symposium paper, that contingent loan plans for repayment of tuition, and their possible defects, can be analyzed along these lines.

REFERENCES

- Arrow, K. J., *Essays in the Theory of Risk Bearing* (Chicago: Markham, 1971).
- , "Higher Education as a Filter," *Journal of Public Economics*, II (July 1973), 193–216.
- Borch, K., *The Economics of Uncertainty* (Princeton, N.J.: Princeton University Press, 1968).
- Hahn, F. H., "Notes on R-S Models of Insurance Markets," mimeo, Cambridge University, 1974.
- Riley, J., "Competitive Signaling," *Journal of Economic Theory*, X (April 1975), 174–86.
- , "Informational Equilibrium," mimeo, Rand Corporation, 1976.
- Revesz, P., *Laws of Large Numbers* (New York: Academic Press, 1960).
- Rothschild, M., and J. E. Stiglitz, "Equilibrium in Competitive Insurance Markets," Technical Report No. 170, IMSSS Stanford University, 1975.
- Spence, M., "Job Market Signaling," this *Journal*, LXXXVII (Aug. 1973), 355–79.
- , *Market Signaling* (Cambridge: Harvard University Press, 1974).
- Stiglitz, J. E., "Perfect and Imperfect Capital Markets," paper presented to Econometric Society Meeting, New Orleans, 1971.
- , "Education as a Screening Device and the Distribution of Income," mimeo, Yale University, 1972.
- , "Demand for Education in Public and Private School Systems," *Journal of Public Economics*, III (Nov. 1974a), 349–86.
- , "Pure Theory of Local Public Goods," in M. Feldstein ed. IEA Conference Volume (Turin, 1974b.).
- , "The Theory of Screening, Education, and Distribution of Income," *American Economic Review*, LXV (June 1975), 283–300.
- Westhoff, Frank H., "The Theory of Local Public Goods," Ph.D. thesis, Yale University, 1974.
- Wilson, Charles A., "Equilibrium in a Class of Self-Selection Models," Ph.D. thesis, University of Rochester, 1976.