

EQUILIBRIUM IN COMPETITIVE INSURANCE MARKETS: AN ESSAY ON THE ECONOMICS OF IMPERFECT INFORMATION* MICHAEL ROTHSCILD AND JOSEPH STIGLITZ INTRODUCTION

Economic theorists traditionally banish discussions of information to footnotes. Without insurance his income in the two states, "accident," "no accident," was $(W, W - d)$; with insurance it is now $(W - a_1, W - d + a_2)$, where $a_2 = a_2 - a_1$. The vector $a = (a_1, a_2)$ completely describes the insurance contract.¹

Demand for Insurance Contracts On an insurance market, insurance contracts (the a '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, $(1) fV(p, W_1, W_2) = (1 - p)U(W_1) + pU(W_2)$, where $U(\cdot)$ represents the utility of money income² and p the probability of an accident. Individual demands may be derived from (1). A contract a is worth $V(p, a) = V(p, W - a_1, W - d + a_2)$. From THE ECONOMICS OF IMPERFECT INFORMATION all the contracts the individual is offered, he chooses the one that maximizes $V(p, a)$. Since he always has the option of buying no insurance, an individual will purchase a contract a only if $V(p, a) \geq V(p, 0) = V(p, W, W - d)$. 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 $V(p, a)$ is quasi-concave. If market odds are given by EF ³ (as they will be if there are relatively few high-risk insurance customers), then y will make a profit. This establishes that a competitive insurance market may have no equilibrium. 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. 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. 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 assumption that individuals know their accident probabilities, while insurance companies do not (which raises more interesting issues), is given in subsection 11.1 below. 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 . The fraction of high-risk customers is X , so the average accident probability is $p = Xp_H + (1 - X)p_L$. Similarly, if individuals are risk-neutral, it never pays to pool; if they are infinitely risk averse with utility functions

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. We assume that companies are risk-neutral, that they are concerned only with expected profits, so that a contract a when sold to an individual who has a probability of incurring an accident of p , is worth

$$7r(p, a) = (1 - p)a_1 - pa_2 = a_1 - p(a_1 + a_2)$$

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. 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 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 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. 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 by any one customer is known to all companies that sell to him. 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. 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 . Othersuch cases can be analyzed, but we trust that the general principle is clear. 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. 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. To illustrate our, mainly graphical, procedure, we first analyze the equilibrium of a competitive insurance market with identical customers.

5 QUARTERLY JOURNAL OF ECONOMICS

the states: no accident, accident, respectively. 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. An equilibrium contract for low-risk types must not be more attractive to high-risk types than a_H ; it must lie on the southeast side of U_H , the high-risk indifference curve through a_H . We leave it to the reader to demonstrate that of all such contracts, the one that low-risk types most prefer is a_{oL} , the contract at the intersection of E_L and U_H in Figure III. 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.

– QUARTERLY JOURNAL OF ECONOMICS

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). 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. It lies above U_L , the low-risk indifference curve through a_L and also above U_H . If y is offered, both low- and high-risk types THE ECONOMICS OF IMPERFECT INFORMATION will purchase it in preference to either a_H or a_L . If it makes a profit when both groups buy it, y will upset the potential equilibrium of (a_H, a_L) . Thus, if contract a is available from a company, so are the contracts $2a$ and $(\frac{1}{2})a$; former pays twice as much benefits (and costs twice as much in premiums) as a ; the latter is half as expensive and provides half as much coverage. 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. 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. Free entry and perfect competition will ensure that policies bought in competitive equilibrium make zero expected profits, so that if a is purchased, $(3) a_L(1-p) - a_2p = 0$. If $ir(p, a) > 0$, then firms offering a lose money, contradicting – THE ECONOMICS OF IMPERFECT INFORMATION

I the definition of equilibrium. However, i_3 offers more consumption in each state than a_H , and high-risk types will prefer it to a_H . If f and a_H are marketed, both high- and low-risk types will purchase A . The nature of imperfect information in this model is that insurance companies are unable to distinguish among their customers. 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. In them we question the behavioral assumptions and the equilibrium concepts used in Section I.

11.1 Information Assumptions

Suppose that there are two groups of customers and that not all individuals within each group have the same accident probability. 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. 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. 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 a_i , in return for which he will be paid W if an accident occurs. It implies, in effect, that the seller of insurance specifies both the prices and quantities of insurance purchased. a^* satisfies the two conditions of equilibrium: (i) it breaks even; (ii) selling any contract preferred to it will bring insurance companies expected losses. 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. 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. The separating equilibrium we have described may not be Pareto optimal even relative to the information that is available.

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. 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.

11.2 Price Competition Versus Quantity Competition

One can imagine our model of the insurance market operating in two distinct modes. 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. 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. Serious consideration of costs of communication, imperfect knowledge, and the like would, it is believed, complicate without informing. 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. 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. Together these assumptions guarantee that any contract that is demanded and that is expected to be profitable will be supplied. 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. 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. Since customers are risk-averse, the point a^* is located at the intersection of the 45-degree line (representing equal income in both states of nature) and the fair-odds line. It follows from (1) that at the slope of the high-risk indifference curve through a , U_H , is $(p_L/(1 - p_L) - p_L)/(1 - p_H/p_H)$ times the slope of UL , the low-risk indifference curve through a . In this figure U_H is a broken line, and UL a solid line. The curves intersect

at a ; thus there is a contract, A in Figure II, near a , which low-risk types prefer to a . The high risk prefer a to A . Since A is near a , it makes a profit when the less risky buy it, $(r(p_L, f)(p_L, a) - r(p, a) = 0)$. The existence of A contradicts the second part of the definition of equilibrium; a cannot be an 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 costs of separating arise from the individual's inability to obtain complete insurance. 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. Our conclusions (or ones very like) must follow from a serious attempt to comprehend the workings of competition with imperfect and asymmetric information. Individuals know that drinking affects accident probabilities, but it affects different people differently. The first, price competition, is familiar to all THE ECONOMICS OF IMPERFECT INFORMATION students of competitive markets. 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. 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. 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. This assumption is defended and modified in subsection 11.1. 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 most competitive markets, sellers determine only price and have no control over the amount their customers buy. 1.5

Equilibrium with Identical Customers Only when customers have different accident probabilities, will insurance companies have imperfect information. The equilibrium policy a^* maximizes the individual's (expected) utility and just breaks even. Purchasing a^* locates the customer at the tangency of the indifference curve with the fair-odds line. If $i - (p, a) > 0$, 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 a , so a cannot be an equilibrium. QUARTERLY JOURNAL OF ECONOMICS This is a^H in Figure III; it must be part of any equilibrium