

Competition and Insurance Twenty Years Later

MICHAEL ROTHSCHILD

Woodrow Wilson School of Public and International Affairs, Princeton University, Princeton, NJ 08544

JOSEPH E. STIGLITZ

The World Bank and Stanford University, Stanford, CA 94305

Abstract

We are honored to address the European Group of Risk and Insurance Economists and will take the opportunity to make some reflections on the rather uneasy relationship between insurance and competition.

Economists generally prescribe competition as a solution for markets that do not work well. Competition allocates resources efficiently and encourages innovation and attention to what customers want. Insurance markets differ from most other markets because in insurance markets competition can destroy the market rather than make it work better.

One of the dimensions along which insurance companies compete is underwriting—trying to ensure that the risks covered are “good” risks or that if a high risk is insured, the premium charged is at least commensurate with the potential cost. The resulting partitioning of risk limits the amount of insurance that potential insurance customers can buy. In the extreme case, such competitive behavior will destroy the insurance market altogether. A simple model illustrates.

Key words: insurance, adverse selection, competitive outcomes

1. A model of underwriting

Suppose that a group of individuals have income $Z = W + \tilde{\eta}$, where $\tilde{\eta}$ is an insurable risk and $E[\tilde{\eta}] = 0$. If there is actuarially fair insurance, risk-averse individuals will purchase insurance against $\tilde{\eta}$ and have income W . Suppose that the risk $\tilde{\eta}$ can be decomposed so that $\tilde{\eta} = \tilde{\varepsilon} + \tilde{\delta}$ where $E[\tilde{\delta} | \tilde{\varepsilon}] = 0$. As a concrete example suppose that $\tilde{\eta}$ represents the risk of getting breast cancer (normalized to have zero mean and converted into monetary units). Then $\tilde{\varepsilon}$ can take one of two values depending on whether the person in question has the breast cancer gene, $BrCa_1$. If the gene is present, $\tilde{\varepsilon} < 0$; if it is not, $\tilde{\varepsilon} > 0$. Here $\tilde{\delta}$ represents the (monetized) risk of incurring breast cancer conditional on one's genetic makeup; clearly, $\tilde{\delta}$ may be dependent on $\tilde{\varepsilon}$. If insurance companies know ε , they will offer policies conditional on $\tilde{\varepsilon}$.

This is a general way of modeling underwriting. When the insurance market divides up risk classes into smaller groups, it is partitioning $\tilde{\eta}$ into the sum $\tilde{\varepsilon} + \tilde{\delta}$ and offering to insure only the conditional risks ($\tilde{\delta} | \tilde{\varepsilon}$) instead of the unconditional risk $\tilde{\eta}$. In the extreme case, $\sigma^2(\tilde{\varepsilon}) = \sigma^2(\tilde{\eta})$, and no insurance is available.

2. Welfare implications

Underwriting decreases expected utility for risk-averse individuals who do not know their $\bar{\epsilon}$. Prior to underwriting, those who bought insurance received a certain income of W . After underwriting, they get the random variable $W + \bar{\epsilon}$. Since $\sigma^2(\bar{\epsilon}) > 0$, risk averters prefer the situation before underwriting.

Arguably this analysis is misleading. Individuals may know what their $\bar{\epsilon}$ is. Those with a good value (those without BrCa₁) will be happy not to be pooled with those who do. This logic drove our previous analysis of the interaction of competition and adverse selection, and we shall return to this point.

The cogency of this argument depends on what one considers the initial situation. Consider a woman who does not know whether she has BrCa₁ but does know that the insurance company will test her before it decides what policy to offer her. For this woman, underwriting clearly decreases expected utility.

The same logic applies if we try to evaluate the distribution of income in a society with and without underwriting. Any concave social welfare function will prefer the distribution of income in a society without underwriting to the distribution of income in a society with underwriting. To see this, imagine yourself behind a Rawlsian veil of ignorance; ask whether you would prefer to be born in a society with underwriting or a society without underwriting. If you are risk-averse, you will choose the society without underwriting. This argument assumes that obtaining information about risks is free. Taking account of the fact that underwriting uses real resources strengthens our conclusion. In many cases risk-averse insurance customers would not get the information that is required to decompose the risk $\bar{\eta}$ into $\bar{\epsilon} + \bar{\delta}$. Competing insurance companies will get this information. Thus, competition works to limit the scope of what insurance will do. In this sense there is a fundamental tension between competition and insurance. One of the natural methods that insurance companies use to compete (underwriting) destroys or considerably limits what insurance markets can do.

Recent and projected developments in genetics¹ make the concern about underwriting of some interest for policy. We are fast accumulating genetic information that will permit tight predictions of future health status. Huntington's Disease is a striking example. Huntington's is a fatal degenerative disease whose symptoms become apparent in young adulthood; Woody Guthrie died of Huntington's. A single dominant gene controls Huntington's Disease. Genetic tests now predict with essentially complete accuracy whether someone will get this fatal degenerative disease. This bids to make it impossible to insure against Huntington's. It would make sense for insurance companies to exclude Huntington's from general coverage and write riders covering Huntington's only to those willing to get tested. However, testing would reveal whether someone would get the disease so there would be no insurance. Adding to the poignancy of this example is the fact that many at risk for Huntington's have decided not to get tested. They have decided they would rather not know whether they will eventually succumb².

Huntington's is a dramatic but not a singular example. Many common diseases (including breast cancer, asthma, and diabetes) have a genetic base. We soon will have the ability to predict with much more accuracy than we now can who will get many diseases. If this

information is available, insurance companies will want to use it for underwriting. The consequence will be a decrease in the risks that people can insure themselves against.

3. Adverse selection

The knowledge that modern genetics will soon be able to provide about people's future health status is dramatic. We now stand behind a Rawlsian veil that is about to be lifted. Policy makers in the United States are trying to grapple with an appropriate response to this coming change. The major policy response in the United States has been to try to prevent insurance companies from using genetic information. This is the recommendation of the prominent group of scientists who work on ELSI (the Ethical, Legal and Social Implications of the Human Genome Project) (see Hudson et al. [1996], p. 391). The Kennedy-Kassebaum bill that just became law in the United States prohibits insurance companies from denying coverage on the basis of preexisting conditions. Most analysts interpret genetic status as a prime example of a preexisting condition. Again many analysts expect that this legislation will be followed by legislation that sharply limits the price differences that insurance companies can charge customers on the basis of preexisting conditions. This, you all understand, has the potential of creating a serious problem of adverse selection.

Adverse selection occurs (in insurance markets)³ when an insurance customer has more information about his or her health than the insurance company does (or can) use⁴. If insurance companies are prohibited from using information about the genetic traits of their customers then a situation in which customers have more information than insurance companies has been created. Adverse selection actually requires only that individuals with different role profiles behave differently in some way which the insurance company can exploit. Differences in information do give rise to such behavioral differences.

Our earlier article studied the desultory consequences of mixing insurance and adverse selection. We review some of that argument here. Before doing so, we think it important to make two points. The first is that even though the maximizing behavior of consumers has some unfortunate consequences, it has beneficial consequences as well. Indeed, most sensible plans to reduce or control the increase of health costs in the United States rely heavily on the assumption that many consumers want to control the amount they spend on health care⁵.

The second is that we believe that the evidence is now good that adverse selection is an important phenomenon. We quote from a leading student of health care in the United States (Cutler [1996], p. 30):

Almost all health insurance systems where individuals are allowed choice of insurance have experienced adverse selection. Medicare enrollees who choose managed care...are healthier than...[those] who do not. The Federal Employees Health Benefits Program...has adverse selection between more and less generous policies. The spread in premiums between more and less generous policies is 68 percent greater than benefits alone would dictate....And almost every large firm that has encouraged employee choice has found the cost of the most generous policies increases sufficiently rapidly that these policies are no longer viable (often termed a "death spiral").

Adverse selection leads insurance customers to pick and choose among different contracts. Our original article analyzed some of the difficulties that arise when companies recognize that consumers with different health risks have different preferences over insurance policies. When companies use this information to compete, there are important consequences. Customers with different health risks will choose different policies. Assume, as we do unless we explicitly state otherwise, that consumers differ only in their health risks. Then, competition ensures that the only possible equilibrium is one in which all the purchasers of an insurance policy have the same risks and one in which each insurance contract breaks even⁶. Competition *separates rather than pools risks* and thus eliminates cross-subsidies among policies.

In our article we noted two further desultory consequences of mixing adverse selection and competition. The first was an observation about social welfare. Consider a *simple* world as one with no moral hazard (the actions of the insured do not influence the probability that he or she will become ill) in which risk-averse consumers purchase insurance contracts from risk-neutral providers of insurance⁷. Elementary arguments imply that in such a simple world with only one class of customers (remember that consumers differ *only* in their riskiness), the best (socially optimal) insurance contract is the one that provides complete insurance at actuarially fair odds.

Now consider a simple world with several classes of customers. Since competition separates, each risk class will purchase a contract that breaks even when only people of that risk class buy it. Call such a set of contracts a *simple separating set of contracts*. Notice that

The set of separating contracts which provides complete insurance to each class of consumers Pareto dominates any other simple separating set of contracts.

The Pareto dominating set of simple separating contracts is a competitive equilibrium when the risk class to which each consumer belongs is publicly observable. (Suppose each person's odds of getting ill were stamped on his or her forehead.) Competitive, risk-neutral companies will not enter into contracts with negative expected value. Insurance companies will refuse to sell to high-risk consumers policies that break even only when purchased by people with low risks.

Now consider what happens when no information about consumers' risk status is available except what insurance companies can glean from their behavior. To insurance companies all potential customers are the same. We showed in our earlier article that firms respond to this dilemma by selling a set of separating contracts. But the set of separating contracts cannot provide complete insurance, for then all customers would attempt to buy the contract that broke even when only those with the lowest risk bought it. But this contract will lose money when everyone buys it and thus cannot possibly be an equilibrium.

It is clear which simple separating set of contracts *can* be an equilibrium. Consider first those with the highest risk. Let them buy the contract—which we denote α_H —that offers complete insurance at their odds. Now consider the group with the next-highest risk. They will prefer the contract that offers complete insurance at their odds to α_H . Among the contracts that break even when the second-highest risk group buys them, consider those that the second-highest risk group finds equal in value to α_H . Select the one of these contracts that maximizes the utility of the second-highest group and call it α_{H-1} . Let

the second-highest risk group buy α_{H-1} . Continue in this way finding break-even contracts that group k finds equal in value to those bought by group $k + 1$, and find the contracts most preferred by group $k + 1$. The sets of contracts $\{\alpha_H, \alpha_{H-1}, \dots, \alpha_1\}$ is the only possible equilibrium with competition and adverse selection.

Note that $\{\alpha_H, \alpha_{H-1}, \dots, \alpha_1\}$ is Pareto dominated by a simple separating set of contracts that provide complete insurance. This welfare-dominant set of contracts would be an equilibrium were there complete information. We thus see that a possible consequence of prohibiting insurance companies from using genetic information is to move from the simple separating set of contracts—that would be an equilibrium were genetic information public—to $\{\alpha_H, \alpha_{H-1}, \dots, \alpha_1\}$.

Our second observation concerned the existence of equilibrium. We established that the only possible equilibrium was the set of contracts $\{\alpha_H, \alpha_{H-1}, \dots, \alpha_1\}$. It is easy, however, to find examples in which $\{\alpha_H, \alpha_{H-1}, \dots, \alpha_1\}$ cannot be an equilibrium. Consider the contract β that offers complete insurance and breaks even when every potential consumer buys it. Those with lowest risk might prefer β to α_n because β offers them more insurance albeit at worse odds than α_n . So might all other groups. Then if a company offered β (or even a contract that offered complete insurance at slightly less favorable odds than β), all consumers would prefer this contract to the set $\{\alpha_H, \alpha_{H-1}, \dots, \alpha_1\}$. This shows that $\{\alpha_H, \alpha_{H-1}, \dots, \alpha_1\}$ may not be an equilibrium; since it was the only possible equilibrium, we concluded that in an insurance market where competition worked through adverse selection, equilibrium might not exist.

Since we wrote our article much hard and good work has been devoted to the question of the existence of equilibrium. Some work has established that the problem is more pervasive than we had indicated. We noted that equilibrium would not exist when there were relatively few high-risk individuals. Then for the low risk the gains from distinguishing themselves from the common pool (being able to buy insurance at slightly better odds) are low relative to the costs of separation (purchasing partial insurance). Riley [1979] established that this kind of argument could be extended to show that equilibrium (of the sort we defined) would not exist in a simple world where agents differed continuously in their riskiness.

Other work has taken issue with the concept of equilibrium that we proposed. We defined equilibrium as a Nash equilibrium in insurance contracts. A set of contracts $\{\alpha_H, \alpha_{H-1}, \dots, \alpha_1\}$ was an equilibrium only if (1) when these contracts were offered and consumers selected from the set to maximize expected utility, each contract made nonnegative profits and (2) there did not exist another contract β such that when the set of contracts $\{\alpha_1, \alpha_2, \dots, \alpha_n, \beta\}$ was offered to consumers and consumers choose contracts to maximize expected utility, β made positive profits.

This seemed to us, in our context, the most natural sort of equilibrium concept. Work since our article has focused on other concepts. We do not present a survey of this work here. Much work that has followed has used (and developed) concepts of game theory to make precise the nature of the game played between insurance companies and customers and focused on the exact sequence of moves that each side (insurers and customers) make. Other work has stressed the calculation of entrants into the markets. As Wilson [1977] pointed out, equilibrium will exist if firms that offer new contracts on the insurance market consider properly what might happen when these contracts drive other contracts out of

the market. Others have considered various kinds of mixed strategy equilibria and shown that they exist. We admire the ingenuity, energy, and technical difficulty of this work, but we persist in our belief that the equilibrium concept we originally proposed is the natural one. We doubt that mixed strategies, reaction equilibria, or even carefully staged games play much role in real insurance markets. In addition, many of the proposed equilibrium concepts seem inappropriate for analyzing a market where the usual assumptions about atomistic competition would seem to describe consumers. Does it make sense for a small firm to worry that its actions will bring about a collapse of the entire market? We leave it to others to judge whether this reaction reflects wisdom or the incipient senility that comes with age and a concern for practical policy.

What then are we to make of our nonexistence result? To us it is an indication that competition does not mix easily with adverse selection and that competitive markets with adverse selection are often unstable. Since the competitive outcome, when it exists, is not a lovely one, this does not strike us as a cause for alarm.

4. Conclusion

An important problem both for theory and policy is to devise ways of providing health insurance that use the power of competition to bring about efficiencies without severely limiting people's ability to insure against ill fortune. As our ability to predict health outcomes improves, this issue becomes more salient. It would be wrong to imply that better knowledge of the genetic basis of disease can bring only disaster. If used properly, better information can make insurance markets work better.

Up to this point we have assumed that individuals can do nothing to affect their health status. In particular, we have assumed that genetic information is, in the sense of decision theory, information of no value. This is clearly false. Indeed, much of the more mundane promise of the genetic revolution lies precisely in the possibility of alerting some of us of the need to take particular kinds of care. Those with a predisposition to high blood pressure and heart attacks have more reason to moderate their consumption of fats than those without such a predisposition. While risk segmentation in general leads to a diminution of welfare, the threat of risk segmentation can lead people to take efficient care.

Similarly, it is possible that science will provide policy makers and insurance companies with good enough objective information about a consumer's health status that risk adjustment can work. Suppose the objective probability of a consumer contracting a given illness were public knowledge. Then risk adjustment would allow markets to operate competitively and provide complete insurance. Insurance companies would compete to sell policies offering complete insurance. The government would reimburse insurance companies for their expected costs—which would be determined by the (publicly known) health propensities of their customers. Such risk adjustment would remove any incentives for companies to try to use their prospective customers' behavior to reveal their health status. To date, policy makers have not had enough good information about health status to make risk adjustment practical. It may be that the revolution in genetics will provide insurers and regulators with enough information to make real risk adjustment possible. If this is so, it will truly be a beneficial revolution.

Acknowledgments

We are grateful for helpful discussions with David Bradford, Martin Hellwig, and Philip Kitcher.

Notes

1. See Kitcher [1996] for a good general introduction to these issues.
2. The standard decision theory approach to the value of information assumes that information per se can have no value. Information is valuable if and only if you can use it to make better decisions than you would in its absence. In this theory the value of information is never negative. Within this framework it is not possible to comprehend people actively not wanting information. Yet this does happen. Letting people avoid unwanted information is an important policy problem that developments in genetics will require us to confront. Economists guided by standard decision theory have little to offer to the study of this and similar problems.
3. The phenomenon is, of course, much more general. Insurance markets are not the only contexts in which one party to a contract has information about the likely profitability of the contract that other parties do not (or cannot) use. The literature is replete with examples and analyses of adverse selection in labor, education, product, and other markets.
4. Laws sometimes prohibit insurance companies from acting on information they may have. In the United States insurance companies cannot offer different contracts to men and women even though their actuarial risks are quite different.
5. Cutler [1996] argues convincingly that the major driver of medical costs is "the over provision of technologically intensive care". He also presents evidence that most expensive technology does not lead to a lasting improvement of health. Typically, the effects of new procedures on mortality disappear after a year. Cutler uses the example of treatment after a heart attack (where angioplasty is now routine) to illustrate. The best and most expensive treatments (for example, angioplasty or open-heart surgery after a heart attack) prolong life—but not for more than a year. (A representative study looks at survival rates from heart attacks as a function of distance from a hospital that can do a new and expensive catheterization procedure; life expectancy for a few months after the attack is negatively correlated with distance; life expectancy for a year after the attack is only slightly related). Arguably this is not a situation that consumers want. If they knew what they spent on health care bought them, they might spend less.
6. We make the usual assumption that in competition firms make zero profits.
7. This assumption can of course be justified in a general equilibrium model of a large economy in which consumers own shares of insurance companies.

References

- CUTLER, D. [1996]: "Public Policy for Health Care," NBER Working Paper 5591. May.
- HUDSON, K.L., et al. [1996]: "Genetic Discrimination and Health Insurance: An Urgent Need for Reform," *Science*, 270, 391.
- KITCHER, P. [1996]: *The Lives to Come*.
- RILEY, J.G. [1979]: "Informational Equilibrium," *Econometrica*, 47, 331–360.
- WILSON, C.A. [1977]: "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory*, 16, 167–207.

Copyright of Geneva Papers on Risk & Insurance - Theory is the property of Kluwer Academic Publishing / Business and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.