International Review of Law and Economics (1989), 9(25–39)

THE ECONOMICS OF COSTLY RISK SORTING IN COMPETITIVE INSURANCE MARKETS

SEVERIN BORENSTEIN

Department of Economics and Institute of Public Policy Studies, The University of Michigan, Ann Arbor, MI 48109-1220, USA

I. INTRODUCTION

In the last twenty years, state and federal policy makers have often been asked to restrict the criteria, or tests, on which insurance companies may base their risk categorization. Insurance companies have responded that it is fair and reasonable to charge groups with different probabilities of death or accident actuarially fair rates for each group. The companies argue that if they are not allowed to make these distinctions, the result will be adverse selection with low-risk individuals self-insuring to a greater extent than is efficient and high-risk individuals overinsuring.

This paper analyzes the efficiency of insurance companies' incentives to use certain types of sorting devices—those that are themselves costly to society. Genetic screening, tests for the HIV (AIDS) virus, and criteria that require extensive monitoring of behavior (e.g., cigarette smoking habits) all use real resources to separate risk groups. Such sorting devices do not necessarily enhance efficiency. Though the decline in adverse selection does increase total surplus, this increase may be more than offset by the resource costs associated with determining individuals' risk levels. Still, the test may be used in private insurance markets, because the gains to low-risk individuals from the sorting come not just from the increase in total surplus but also include a transfer of wealth from high-risk people. So long as the low-risk people would be made better off by distinguishing themselves, they will do so, and competitive firms will find it profitable to offer them lower insurance rates.

The analysis is in the tradition of Spence's (1973) signalling theory.[1] Just as education yields a greater private return than social return in Spence's model, the next section demonstrates that risk-sorting yields a greater return to low-risk individuals than to society as a whole. The basic point of inefficient costly risk-sorting has been made by Crocker and Snow (1986) who have also shown that price/quantity bundling to induce self-selective sorting, as proposed by Rothschild and Stiglitz (1976), may be imposed inefficiently. Section II gives some intuition for these results.

Sections III and IV extend the Crocker and Snow analysis in two directions. Section III considers the role of a consumer's information about his own risk

I thank Ted Bergstrom, Paul Courant, David Lam, Jeffrey Mackie-Mason, Dan Rubinfeld, Hal Varian, and two anonymous referees for their helpful comments. Bruce Schillo provided excellent research assistance. Any remaining errors are my own.

[1]Since Spence's seminal work, an extensive literature on inefficient self-sorting in labor markets has evolved. See Arrow (1973), Stiglitz (1975), Salop and Salop (1976), Wolpin (1977), and Weiss (1983). A parallel literature has developed on self-selective price discrimination. See Salop (1976), Varian (1981), and Borenstein (1985).

level. Even if consumers are unsure of their own riskiness, it is shown that any information advantage that consumers may have over insurance companies can lead to inefficient, costly sorting. Section IV analyzes complications from moral hazard. When consumers can influence their risk levels, the possible inefficiency remains. In that case, the extent of the problem is a function of the correlation in the population between preference for risky behavior and preference for insurance.

Section V discusses the implications of this work for the current controversy regarding insurance testing for the HIV (AIDS) virus. The theoretical results in earlier sections are shown to imply that HIV testing is likely to be used more frequently than is efficient. Section VI concludes the paper and discusses some further extensions of this work.

II. INEFFICIENT SORTING IN A SIMPLE MODEL OF COMPETITIVE INSURANCE MARKETS

Consumers in this market have either a low probability of an "accident," L, or a high probability, H. Each person knows the risk category to which he or she belongs, but there is nothing that can be done to change groups. (Medical history or driving record would be examples of such a distinction.) The proportion of the population that is high risk is β, where $0 < \beta < 1$. People within each risk classification are identical, though low-risk individuals may have different preferences than high-risk people. Insurance companies are risk neutral and all consumers are risk averse. Insurance companies sell contingent contracts that pay \$1 in the event that the insured person has an accident. We assume that contracts are sold under constant returns to scale and that all companies earn zero economic profits.

Inelastic Demand for Insurance

We consider first the case of infinitely risk-averse consumers who, thus, will always fully insure. That is, each person has a completely inelastic demand for insurance (contingent contracts) at the quantity equal to the loss that would occur if an accident took place. Perfectly inelastic demand, of course, insures that there is no adverse selection and, thus, no efficiency loss from pooling different risk types. If all high-risk people buy D_H and all low risk people buy D_L, then the pooled price would be $P = D_L(1 - \beta)L + D_H\beta H$.

Though there is no gain in total welfare from risk sorting in this case, a low-risk individual still stands to gain $D_L(P - L)$, or area A in Figure 1, if she can prove she is low risk. Thus, a test will be imposed under competition if the per person cost of the test, T, is less than the subsidy that each member of the low-risk group would pay with pooling, area A. Essentially, low-risk individuals are willing to pay T in order to receive an income transfer of A from the high-risk people so long as $T < A$. If T were only slightly less that A, the test would still be used even though the loss to high-risk types would dwarf the gain to low-risk types.

Non-Zero Demand Elasticity for Insurance

With utility functions that are strictly concave in income, consumers of each type will have a downward sloping demand for contingent insurance contracts. Figure 2 presents the demand curve for a member of each risk classification. The quantity demanded by a high-risk person is drawn to be greater at any price than the quantity demanded by a low-risk person. This seems reasonable, because the

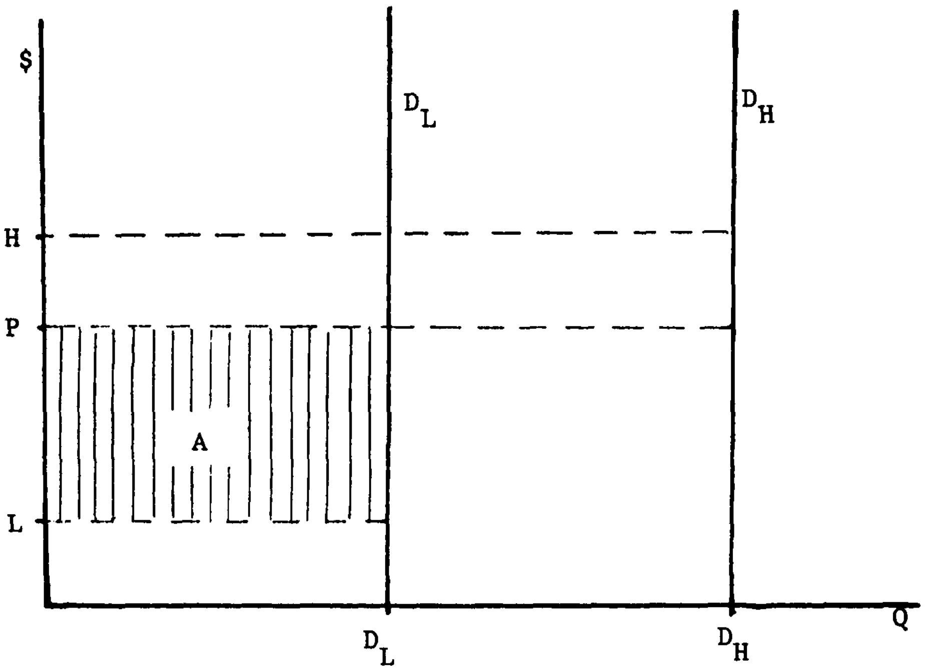

FIGURE 1. Pooling and sorting with zero demand elasticity

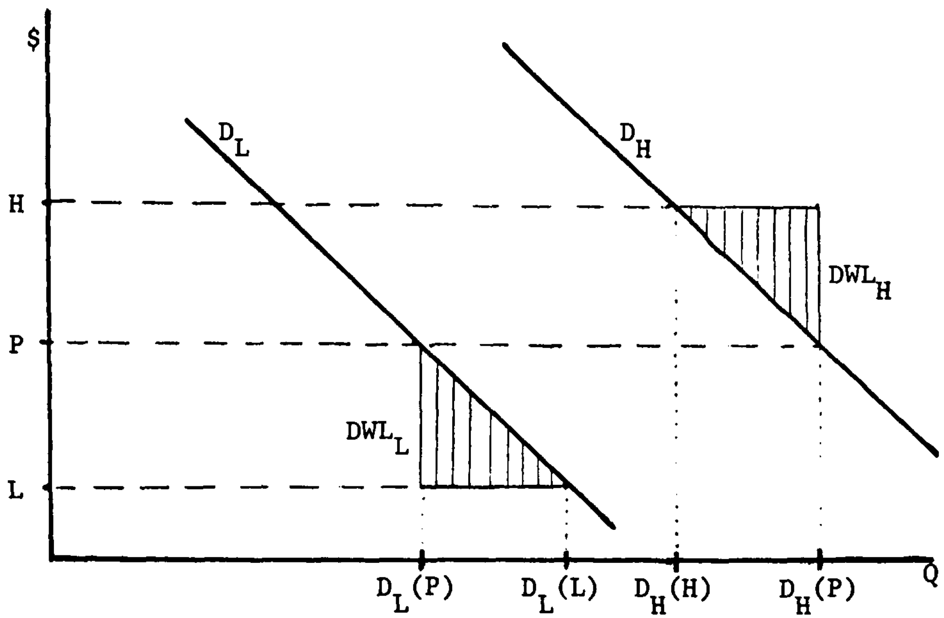

FIGURE 2. Pooling and sorting with non-zero demand elasticity

expected payoff is larger for a person at greater risk, but relative magnitudes of demand do not affect the analysis.

Deadweight loss is minimized at zero in this market when all low-risk buyers purchase $D_L(L)$ one-dollar contracts and all high-risk consumers purchase $D_H(H)$.[2] This will occur in a competitive insurance market if members of each type can be costlessly distinguished. If insurance companies do not sort consumers by risk level, pooling will result, with a price, P, defined implicitly by

$$P = \frac{(1 - \beta) * D_L(P) * L + \beta * D_H(P) * H}{(1 - \beta) * D_L(P) + \beta * D_H(P)}. \tag{1}$$

Since P is a weighted average of L and H with strictly positive weights, it follows that $L < P < H$. Thus, if a pooling equilibrium obtains, there will be inefficiencies from the overinsuring of high-risk individuals (DWL_H) and the underinsuring of low-risk individuals (DWL_L), the standard adverse selection problem.

We now again suppose that there is a way of testing individuals, at a cost of T per individual to determine whether or not they are low risk.[3] We assume that the test is optional—a person can forego it and be classified as high risk—and that the test is errorless.[4] Under these conditions, only low-risk people will take the test.[5]

The condition under which it is welfare improving to use the test is straightforward.[6] The test should be imposed if and only if

$$(1 - \beta) * T < (1 - \beta) * DWL_L + \beta * DWL_H. \tag{2}$$

The deadweight loss due to adverse selection must be greater than the resource cost of imposing the test. Alternatively, the compensation principle must be satisifed, which will be the case if the dollar-valued gains to the beneficiary group

[2]If the utility functions are not state dependent and the financial loss to a person from an accident is still the same whether he is low or high risk, then $D_H(H) = D_L(L)$. Faced with actuarially fair insurance, a risk-averse person will fully insure.

[3]It is assumed that the cost will be paid by the consumer, but in a competitive insurance market, this is not important to the results.

[4]A requirement that all individuals be tested will only strengthen the results developed below.

[5]The test need not even be imposed 100% of the time for it to deter high-risk participation. If an individual tested must pay for the test whether he passes or not, then high-risk people will be deterred if the probability of being tested, pr, is high enough so that

$$pr * (CS_H(H) - T) + (1 - pr) * (CS_H(L) - T) < CS_H(H),$$

where $CS(\cdot)$ is consumer surplus as a function of price. In fact, by charging an infinitely large fee to those taking the test and refunding all but T only to those who pass, insurance companies could discourage high-risk individuals from claiming to be low risk while imposing the test an arbitrarily small proportion of the time. The infinite penalty to those who fail is, of course, not realistic, but this does point out a procedure that could lessen the inefficiency. Still, the discussion that follows is unchanged in principle if the test is imposed with some probability between 0 and 1 so long as the probability is high enough to deter high-risk types.

[6]The discussion of welfare here refers to deadweight loss or total surplus in this market. These are inexact measures of the global efficiency of the insurance prices, but the results are unchanged in the more general, precise analysis, as shown in footnote 9.

are large enough to fully compensate for the losses incurred by others, that is, if the net change in surplus in the insurance market is greater than the cost of testing.

This is not, however, the decision criterion that would be used in a competitive market. If the gain to low-risk buyers from imposing the test is greater than the cost they pay, competition will drive firms to offer the test and then charge a per contract price of L to the people who are shown to be low risk. Rather than comparing the cost of the test to the decrease in deadweight loss or change in consumer surplus of both groups, a competitive market would compare the test cost to the change in consumer surplus of only a low-risk individual.[7] The market comparison omits consideration of the loss imposed on high-risk individuals. Thus, one might find the market using a costly test that increases the welfare of low-risk buyers only slightly after they pay the cost of the test, while it decreases the welfare of the high-risk group by a much greater amount.[8]

Result 1: In a competitive insurance industry, companies may use some costly sorting devices that decrese total surplus.[9]

Though the pooling equilibrium may result in less deadweight loss than a competitive testing equilibrium, pooling does not Pareto dominate the testing equilibrium. The loss of consumer surplus to the high-risk group from testing can be greater than the gain to the low-risk group, but only imposition of the costly test will induce high-risk individuals to reveal themselves. If a scheme were worked out for compensation of low-risk individuals in lieu of testing, any individual would be better off claiming to be low risk.

[7]Consumer surplus is equal to total surplus in this case, because insurers earn no profits.

[8]This argument can also be made in terms of deadweight loss. The test would be used if, and only if,

$$T < (D_L(P) * (P - L) + DWL_L).$$

It is clear that if

$$(1 - \beta) * D_L(P) * (P - L) > \beta * DWL_H,$$

then the competitive criterion permits use of some tests that increase deadweight loss. In a competitive market, however, this condition always holds. To see this, note that $D_L(P) * (P - L)$ is the payment that each low-risk person makes above the cost of insuring him when there is pooling. Since P is defined to be the breakeven price under pooling, the total subsidy *from* all low-risk people must equal the total subsidy *to* all high-risk people. That is,

$$(1 - \beta) * D_L(P) * (P - L) = \beta * D_H(P) * (H - P)$$

From figure 2, it is apparent that the term on the right includes and is strictly greater than $\beta * DWL_H$.

[9]The result can be shown more briefly, though less intuitively, in a more general model. In this analysis, the compensation principle implies

$$(1 - \beta) * [Y_L - E_L(V_L(Y_L,P),L) - T] > \beta * [E_H(V_H(Y_H,P),H) - Y_H],$$

where $E(\cdot)$ is an expenditure function and $V(\cdot)$ is an indirect utility function. The left-hand side is the gain to the low-risk people from the test (presumably positive if the test would ever be used). The right-hand side is the loss (also positive) to the high-risk people. The result is shown by noting that the test will be used so long as the left-hand side is greater than zero, a weaker test than the compensation principle demands.

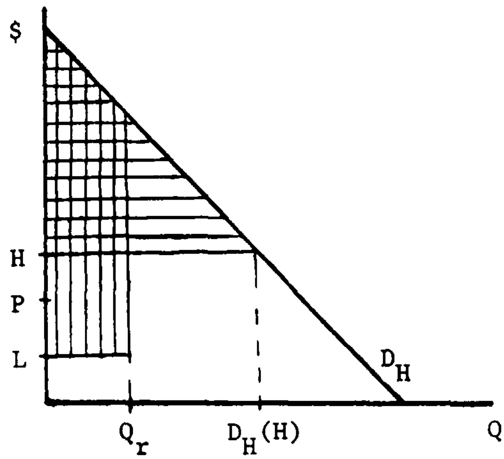

FIGURE 3. Choice facing high-risk individuals under self-selection

Inefficient Self-Selective Sorting

Rothschild and Stiglitz (1976) examined an alternate method of sorting that does not require an explicit test. They pointed out that adverse selection may be overcome with self-selective sorting by offering price/quantity bundles of insurance where lower prices are associated with smaller quantities of insurance and, thus, with a higher level of self-insurance on the part of the buyer.[10] They showed that with such bundles, the only equilibrium that can exist is one in which only low-risk individuals are induced to buy the low-priced insurance. The high-risk individuals are separated out, because self-insuring is more costly for them than for low-risk types.[11]

Figures 3 and 4 demonstrate the conditions under which a self-selection equilibrium will exist. If the quantity of insurance that a person can buy at the low price is restricted to Q_r, an amount less than high-risk people would buy at the separating price, $D_H(H)$, then high-risk individuals may choose to buy all the insurance that they want at price H rather than buy a limited quantity at price L. Writing consumer surplus as a function of both price and quantity now, because rationing may cause individuals to be off their demand curves, the first condition for a self-selective equilbrium to exist is that

$$CS_H(H,D_H(H)) > CS_H(L,Q_r), \tag{3a}$$

that is, high-risk people prefer to fully insure at price H than to buy Q_r at price L.[12] The horizontally striped area in Figure 3 must be larger than the vertically

[10]Insurance companies offer such bundles by giving lower prices on policies with higher deductibles. Their motivation, however, may have as much to do with moral hazard problems as with adverse selection.

[11]Implicit in the analysis is the assumption that high-risk individuals get greater surplus from any given quantity of insurance than low-risk individuals. This is satisfied if the two types do not differ in their risk aversion or the loss they suffer if an accident occurs.

[12]As with the analysis in the previous subsection, the change in total surplus is an imprecise measure of overall welfare change.

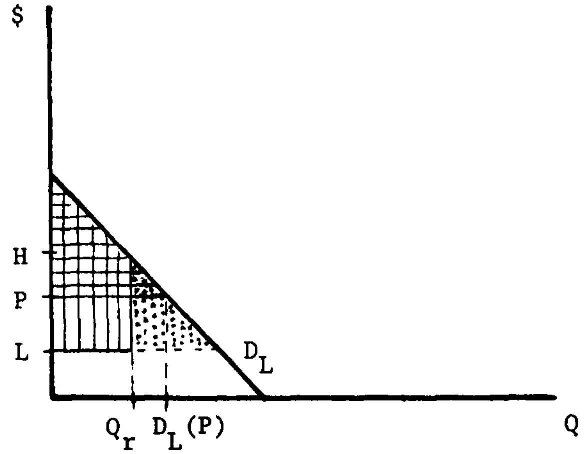

FIGURE 4. Choice facing low-risk individuals under self-selection

striped area. At the same time, however, low-risk individuals must be better off with the low-price, restricted-quantity bundle than under pooling:

$$CS_L(L,Q_r) > CS_L(P,D_L(P)), \tag{3b}$$

where P is from Equation 1. Otherwise, a different firm could offer unrestricted quantity at price P and attract both low- and high-risk types. The horizontally striped area in Figure 4 must be smaller than the vertically striped area.

As Q_r decreases from $D_L(L)$, (3a) is more likely to be satisfied, but (3b) is less likely to hold. Depending on differences in size, risk level, and risk aversion of the two groups, and on the size of the loss, Z, there may or may not be a Q_r for which (3a) and (3b) hold simultaneously. If such a separating equilibrium does exist, it too could obtain even if it results in lower total consumer surplus than pooling.[13] The cost of this type of sorting (or test) is the loss in expected utility to the low-risk group versus costless separation, because they cannot fully insure in this case. In Figure 4, this cost is the loss in consumer surplus represented by the dotted area. If the equilibrium Q_r leaves low-risk people only slightly (or infinitessimally) better off than under pooling, then the loss to high-risk types

[13]Rothschild and Stiglitz also show that pooling cannot be an equilibrium if firms are permitted to offer these price/quantity bundles. To see this in the terms of figure 3, note that the Q_r for which $CS_H(L,Q_r) = CS_H(P,D_H(P))$ necessarily has $CS_L(L,Q_r) > CS_L(P,D_L(P))$. Stated differently, start from a bundle with price equal to L and the associated restricted quantity, Q_r, equal to $D_L(L)$. If this would not yield separation, begin reducing Q_r and note that such a change will decrease consumer surplus more rapidly for the high-risk person consuming this bundle than for the low-risk person. Wilson (1977) shows that a self-selective equilibrium will always exist if the firms have somewhat more forward-looking conjectures than in Rothschild and Stiglitz. This is because a firm that might break the separating equilibrium with a pooling contract would recognize that it would in turn be vulnerable to losses when a new separating contract skims off only its low-risk customers.

from separation will be greater than the gains to low-risk types.[14] Nonetheless, so long as there is some Q_r that yields an equilibrium increase in the consumer surplus of low-risk individuals versus pooling, self-selective sorting will obtain, regardless of the resulting loss in surplus to the high-risk group.

III. THE ROLE OF CONSUMERS' INFORMATION

Inefficient use of some sorting devices can result when consumers who know they are low risk insist on taking the costly test and then receiving the price L. In many cases, however, consumers do not know whether they are high risk or low risk. For example, an individual is unlikely to know *ex ante* whether a test for high serum cholesterol would indicate that she has a relatively high risk of developing heart disease.

An extreme case of poor consumer information occurs when every consumer's best guess of her probability of passing a test is the proportion of the entire population that passes. That is, the consumer has no better idea of her chance of passing the test than the insurance company has. In that case, with buyers and sellers of insurance possessing the same information, there is no adverse selection problem.

With symmetric *ex ante* information of this sort, no buyer would pay for a test that tells her whether she is high or low risk if the insurance company necessarily also received that information. To see this, note that if the buyer and the competitive insurance company have the same information, then insurance will be offered at what the buyer perceives to be an actuarially fair price. In response, the risk-averse buyer will fully insure. A test will change the price a buyer faces, but since she will still perceive it to be actuarially fair, she will still fully insure against an accident. Thus, a test will not change the quantity of insurance that an individual will buy. Since the companies break even before and after the test, while selling the same quantity to each consumer, the pretest price that all consumers face must be equal to their expected posttest price. The test would increase the *ex ante* variance of insurance premiums without decreasing any consumer's expected payment. Thus, in this extreme case, no costly sorting devices would be used.[15]

Still, if some consumers think that they have a higher than average probability of passing the test, they might, depending on their degree of risk-aversion and the cost of testing, perceive an increase in their expected utility from taking the test. The decline in expected utility that others in the population suffer from the test would not be considered in a competitive market. If the lower-risk consumers believe that their probability of passing is only slightly higher than the average, then the difference between socially optimal and private incentives may be small.

[14]Hoy (1982) points out the "knife-edge" Pareto comparison that is possible. As the gain to low-risk types from separation approaches zero, pooling Pareto dominates the separating equilibrium.

[15]The point is related to one made by Chamberlin (1985). He argues that many tests would be banned if people had to make the choice regarding their legality from a Rawlsian "original position," i.e., not knowing which category they would fall into under the test. So long as the adverse selection problem (once people learned their status) were not too great, the lower variance in insurance rates without the test (as seen from the original position) would be more valuable than the efficiency gain from using the test. Stated differently, people would like to be insured not just against an accident but also against having or developing the characteristics that would put them in a high-risk category. Also, see Schmalensee (1984) for an analysis of the equitability of imperfect sorting.

Still, the difference will exist so long as at least one individual believes that her probability of passing the test is higher than average.

If uncertain consumers can take a test without necessarily revealing the results to their insurance companies, the possibility of inefficient, costly sorting increases. Such an option would have no effect on whether costly testing would increase total surplus, but "optional-disclosure testing" would increase the private incentive to get tested, provided that a finding of low-risk could then still be used to obtain lower insurance rates. If the test results are automatically revealed to the insurance company, then the test will be used if any consumer believes that

$$\pi * CS(L) + (1 - \pi) * CS(H) - T > CS(P) \tag{4}$$

where π is the consumer's subjective probability of passing the test and $CS(\cdot)$ is her consumer surplus as a function of price.[16] With "optional-disclosure testing," failing the test will result only in remaining in the pooled group. The comparison then would be

$$\pi * CS(L) + (1 - \pi) * CS(P) - T > CS(P) \tag{5}$$

which is a weaker criterion since $P < H$. Clearly, making such tests unavailable in order to avoid asymmetric information is not feasible or, probably, constitutional. Still, the possibility of inefficient use of optional-disclosure tests is eliminated if the test results cannot then be used in rate setting. As discussed above, the cost of such a policy would be adverse selection.

IV. THE EFFECT OF MORAL HAZARD

Moral hazard is introduced into the model by assuming that people can choose whether to be low risk or high risk. This can be interpreted as the consumer's choice of whether or not to engage in a risky activity or, conversely, as his choice of the level of precaution to take against an accident. The model is further generalized by allowing individuals to have different utility functions. In addition to income and the price of insurance, an individual's choice of low- or high-risk behavior now also affects utility.

There may be some people who will choose to be low risk even when there is no testing. These people, for instance, do not like cigarettes. They would not smoke even if doing so would have no effect on their life insurance rates. Such people, designated type I, satisfy

$$V_{L_i}(Y_i,P) > V_{H_i}(Y_i,P). \tag{6}$$

where $V(\cdot)$ is (indirect) utility as a function of income and the price of insurance, L and H indicate different indirect utility functions when the person is choosing to be low or high risk, and i indexes the consumer.[17] For all others, type II individuals, (6) is not satisfied, so they choose to be high risk under pooling.

It is useful to consider first a population in which no one would choose to be

[16]The comparisons in (4) and (5) use consumer surplus, an inexact measure, as was done in section II. The more general statement, analogous to that in footnote 9, is omitted for brevity but is equally straightforward.

[17]The fact that P is determined endogenously does not change the usefulness of these distinctions. For a given pooling equilibrium with price P, the population can be partitioned according to those for whom (6) is and is not satisfied.

low risk—nobody would take safety precautions or avoid high-risk behavior—unless he were rewarded with lower insurance rates. If no separation (and thus no such reward) based on risk level were possible or allowed, the resulting pooled price, with all customers engaging in high-risk behavior, would be H. In that case, no cross-subsidy would exist within the pooled group. If some people then removed themselves from the high-risk pool by becoming low risk and passing a test to prove it, the competitive price to the people who remained in the high-risk pool would still be H. Thus, the competitive criterion for use of the test is consistent with the compensation principle in this case; the change in the individual's welfare from deciding to move to the low-risk group would constitute the entire change in total welfare from the decision.

Result 2: When there are no people who would choose to be low risk under pooling—no type I *individuals—the competitive criterion for use of a costly test is consistent with the compensation principle.*

Result 2 highlights the relationship between the change in price to the nontested group, when the market goes from pooling to separation, and the optimality of the competitive criterion for offering the test. If the price to the nontested group when the test is used—call this price N—is the same as the price without testing, the competitive criterion is optimal. It is more likely, however, that the nontested price would differ from the pooled price. If $N > P$, then use of the test makes nontested individuals worse off. This decline in welfare for the nontested group, however, does not affect the competitive incentive to offer the test to those who would benefit from taking it.

Whether or not introduction of the test raises the price to those who do not take it will depend upon the mix of individuals who choose to take the test (and to be low risk). If, for instance, there were a positive population of type I people and they all took the test when it became available, then only people choosing high-risk behavior—some of the type II's—would remain in the nontested group and the price to that group would be $N = H$.[18] This would be above the price that obtained with pooling, when the type I people pulled down the competitive price to the pooled group, $P < H$.

Though the $N > P$ outcome seems the most likely, the opposite result is also possible. Testing could lower the price to the nontested group as well as to the tested group, resulting in $L < N < P$. This could obtain if, for instance, type I people had systematically smaller demands for insurance than type II's so that for many of the type I consumers, it would not be worth the fixed payment of T to receive the low per contract price for insurance. The mix of people in the nontested group could then contain a higher proportion of people choosing low-risk behavior than had the pooled group when no tests were used. More precisely, the proportion of \$1 contracts purchased in the nontested group by people behaving in a low-risk way could be higher than the proportion they purchased in

[18] Some type II individuals who choose low-risk behavior with availability of the test may still be worse off than under pooling. It may be the case that

$$V_{H_i}(Y_i,P) > V_{L_i}(Y_i,L), \qquad \text{but} \qquad V_{H_i}(Y_i,N) < V_{L_i}(Y_i,L)$$

for some individuals. These people would prefer pooling and high-risk behavior, but when availability of the test causes the nontested price to rise, they prefer to switch to low risk and pay L. In a model with many risk types, the availability of the test could cause a complete "unraveling" in which all people choose to switch to lower risk with the test, but only a few, the equivalent of type I's, prefer that to the pooling equilibrium.

the pooled group when no test was used. The competitive price to the nontested group would then be lower than the price that had been charged to the pooled group in the absence of testing. These low-risk people would remain in the nontested group, while many formerly high-risk people might switch to low-risk behavior and take the test. In this case, the availability of testing would increase the welfare of even those who were not tested. This positive effect on the nontested group would be ignored in a competitive market and, thus, some tests that would increase total surplus may not be offered.[19]

Still, it is difficult to see how this result would be obtained very often. Type *II* people get positive utility from being high risk, at least when insurance sells for P, while type *I* people do not. Type *II* people would have to pay T *and* give up their high-risk activity in order to receive price L, while the cost to type *I* people would be only the price of the test, T. Thus, it seems quite likely that a higher proportion of the type *I* than the type *II* people would take advantage of the testing option if it were available. The result in that case would be $N > P$.[20]

Result 3: If the proportion of contingent contracts purchased by type I *individuals in the nontested group is less than the proportion purchased by type* I *individuals under pooling, then* N > P *and the competitive criterion for using the test is less strict than the compensation principle would dictate (that is, 'too much testing' under competition). If the proportion of contingent contracts purchased by type* I *individuals in the nontested group is greater than the proportion purchased by type* I *individuals under pooling, then* N < P *and the competitive criterion is more strict than the compensation principle would dictate (that is, 'too little testing' under competition).*

V. TESTING FOR THE AIDS VIRUS

As with many ailments, modern medicine has developed a reliable test for increased risk of AIDS before it has found a cure for the disease.[21] The highly accurate Western Blot test for antibodies to the HIV virus that causes AIDS costs about \$50 per individual. A less conclusive, but much less expensive test, called ELISA, is available for about \$10. Because an ELISA test will yield false positive results in 1% of all uninfected cases, a standard ELISA–ELISA–Western Blot

[19]This result is not quite as strange as it might at first appear. Imagine a case in which nonsmokers buy very little insurance because they are quite physically fit, have a healthy diet, and otherwise live a clean (and long) life. It is not worth it for them to show that they are nonsmokers, since the cost of demonstrating this fact is large in comparison to their total insurance payment. Still, under pooling, they are burdened with high prices due the cost of subsidizing smokers, who live unhealthy lives and buy much more insurance. Even though the nonsmokers (under pooling) would all stay in the untested group when a test became available, the offering of a test might still benefit them by inducing some smokers to quit smoking, leave the untested group, and obtain lower insurance rates. For the counterintuitive result to obtain, these ex-smokers must find it more worthwhile to take the test than the nonsmokers, perhaps because they still take part in other unhealthy, and undetectable, behavior that raises their demand for insurance (or because ex-smokers are, for some reason, more risk-averse than nonsmokers). If this were the case, the result would be that the people who were nonsmokers all along would then be burdened by a smaller proportion of smokers in the insurance group (nontested) to which they belong.

[20]As explained in the previous paragraph, the precise statement of this comparison is in terms of shares of contracts purchased rather than shares of individuals of each type in each group.

[21]Costly tests now also exist for the genetic markers that indicate predispositions towards heart disease and Huntington's chorea.

testing sequence has been established to minimize costs of detection and still control false positives. If the first ELISA test is positive, a second one is run. If that is also positive, the Western Blot test, with a false positive rate of 0.5%, follows. HIV infection is inferred if all three tests are positive. The sequence is stopped if a negative result obtains at any point.[22]

The average cost of the testing sequence depends on the infection rate in the population being tested. If most people were infected, then the full sequence would be run on most people and the average testing cost per person would be nearly $70. If very few were infected, then even with the false positives that ELISA occasionally yields, most people would be labeled as seronegative on the first round. In that case, the per person cost would average closer to $10.[23]

The cost of the HIV test is rather small, but so is the probability that any one individual will test positive. A widespread campaign of testing could be quite costly in comparison to the deadweight loss that it would prevent. Moreover, the test might have to be administered periodically because infection can occur or become detectable at any time.[24] Finally, the test costs generally estimated do not include the time cost to the individual being tested, which could amount to a substantial portion of the total cost. Efficiency in the use of these tests would be improved if they were used selectively, with the prescreening based on sexual preference, race, or nationality. For ethical and public policy reasons, however, such distinctions are unacceptable.

If the adverse selection problem is not great and the number of infected individuals is a small proportion of the population, the HIV test may not be worth using even for those who know they are seronegative. That is, the cost of the test could outweigh the benefits to low-risk individuals, so private insurance markets would choose not to use the test. This is more likely to be the case the more limited is the exposure of the insurance company if it (unknowingly) insures seropositive individuals. Thus, for smaller life insurance policies and for medical coverage, the tests may be omitted or optional.[25] Less extensive use of the test is also more likely if the psychological costs of testing are quite high even for those who are extremely unlikely to test positive. On the other hand, in cases where the opportunity for extreme adverse selection is greater, such as with life insurance policies in the millions of dollars, both the efficiency argument for testing and the private incentive to test will be much greater.

It is quite possible that the decrease in rates on some policies (compared to pooling) for those who test seronegative would be greater than the cost of the test, but not so much greater as to satisfy the compensation principle. Those who

[22]The false positive rates cited here are "Best-Case" results when tests are performed under ideal and monitored conditions. In actual labs, the rates may be much higher. Because both the Western Blot and ELISA tests indicate only the presence of antibodies to the HIV virus, which may not appear for one to six months after infection, either test will return false negative results for some period of time following infection. See Miike (1987) and Meyer and Pauker (1987).

[23]These cost estimates are *marginal* cost, given that blood chemistry is already being analyzed for other reasons. If blood tests are to be run specifically for HIV testing, the average cost is likely to be around $75 for a low-infection population. See Mast (1987).

[24]This is not the case for most individual health or life insurance, because the policies are commitments by the insurer to cover diseases contracted after insurance begins. Group policy rates, however, are generally experience-rated and are updated periodically. If a group, *e.g.*, a company, were to have an abnormally high rate of AIDS cases, it might want to have all other members tested to show that the cases found were aberrations.

[25]Estimates of the medical expenses from the time of AIDS diagnosis to death are in the range of $50,000. See Clifford and Iuculano (1987).

test positive or refuse to be tested would face greatly increased insurance rates or would become uninsurable. Furthermore, the costs to those found to be seropositive could extend far beyond the insurance market—to discrimination in jobs, housing, and social settings—if strict confidentiality were breached.

Within the insurance industry the debate is not over *whether* HIV tests should be used but how extensively they should be used. At what size policy does it become cost-effective to require the HIV antibody tests? The calculation done within the industry compares the cost of the test (which is usually paid by the insurer) to the expected savings from avoiding writing policies for HIV-infected individuals.[26] When the decrease in expected payout on a policy—which in a competitive market translates into a decrease in premiums for those who pass the test—is greater than the cost of the test, the insurance company decides, or is forced by competition, to use the test.

Ignored in this calculation, of course, is the loss to the small segment of the population who fail the test and become uninsurable or insurable only at much higher rates. It is clear then that the margin at which the insurance industry will draw the nontesting limit—the limit on the size of policy that may be purchased without submitting to the HIV test—is likely to be inefficiently low. If a test is just barely cost-effective, that is, just barely gives an *ex ante* expected net benefit to those who are most likely to pass the test, then inclusion in the calculation of the negative impact on those who would fail the test will certainly make the test inefficient on total surplus grounds.

This sort of marginal decision-making regarding which policies will require a costly test is common in the insurance industry. For decades, insurance companies have had "nonmedical limits" below which policies would be written without a medical exam. The decision criterion that companies use to set these limits is analgous to the competitive criterion described in Section II. Thus, the inefficiency due to excessively low limits is probably present in many other sorting criteria.

In forming a public policy on HIV testing for insurance, the decision maker should be aware of the negative externalities that a private market decision to test would impose. For some policies, these effects could outweigh the possible gains from testing, particularly if adverse selection is not a severe problem in the market. Of course, there may also be positive externalities from insurance testing for the HIV virus, such as greater awareness of the dangers and of one's own viral status, as well as counseling on methods of reducing the risk of acquiring and transmitting AIDS. If mandatory HIV testing for medical or life insurance would increase awareness and discussion—a view that is controversial among public health experts—then these positive externalities would also have to be considered.[27]

[26]For instance, assume that a population has an overall infection rate of 1% and that the average cost of testing with this infection rate is \$15. If the expected present cost of insuring an HIV-infected person is \$3 per \$1000 of insurance higher than the expected present cost of insuring an uninfected person, then the test becomes cost effective for policies of greater than \$5000. In fact, calculations of this sort yield results that suggest much lower nontesting limits than are actually used by most companies. See Brodrick and Beal (1988). This may be explained by the nonmonetary costs that the test imposes on even those with low probability of infection. Such nonmonetary costs borne by uninfected consumers due to the test would cause the competitive market to raise the nontesting limit, as appears to be the case.

[27]See Barry, Cleary, and Fineberg (1986), Clifford and Iuculano (1987), and Schatz (1987) for differing views on the information effect of insurance testing.

VI. CONCLUSION

This paper has explained why competitive insurance markets do not necessarily make efficient use of costly information. The basis for the result is simply that some people are made worse off when low- and high-risk people are put in separate risk categories by a costly test. This decline in the welfare of those who do not take and pass the test is not considered when a competitive company offers a low rate to people who can show that they are low risk. Because this loss is ignored, tests may be used that benefit a low-risk group only slightly and harm the remaining customers much more.

It is worth contrasting this result with the efficiency of costless sorting. As with costly sorting, costless separation of risk groups harms those who are put in the high-risk group. In that case, however, Crocker and Snow have shown that the losses to high-risk individuals are necessarily smaller than the gains to low-risk individuals. In contrast, the gains to low-risk people from costly sorting may be largely expended in the sorting procedure and, thus, their net gains may be arbitrarily small.

The work of Crocker and Snow has been extended here to analyze cases in which the consumer is unsure of her own risk level and to cases in which risk level is under the control of the consumer. The potential for inefficient risk-sorting generally persists in both cases, though the result may disappear or even reverse in some special circumstances that are identified.

The simplicity of the basic result allows one to see easily the course that extensions to other cases would follow. For example, the test used might have a positive error rate. Nonetheless, if the low-risk types received a low enough rate upon passing the test to justify both the cost of the test and the increased price variance (*ex ante* uncertainty), the negative impact on other consumers, whose expected price would be likely to rise, would still be ignored. The incentive to use some inefficient tests would remain.

The basic result is also robust to the number of risk types. Again, there would in all likelihood be some group of people who would pay a higher price because of the test. If such a group existed, their disutility would be ignored in a competitive market. Any person who is made better off by paying T and receiving a (perhaps expected) lower price would be accommodated regardless of the effect on others buying insurance. Thus, the assumptions of an errorless test and a discrete number of risk categories are not important in determining the inefficiency result.

The results presented here indicate that social welfare might be improved by the prohibition of certain sorting devices used in insurance risk classification. Strictly interpreted, these results do not apply to such low-cost distinctions as sex, race, and age. A broader interpretation, however, might include recognition that when a group that has historically suffered discrimination is again separated out in insurance rate making, people in the group lose more than simply the higher rates that they might have to pay. The gains to the low-cost class might be less than the total loss to the high-cost class if the distinction also reinforces negative stereotypes and discrimination against the latter group.

REFERENCES

Arrow, K. J., "Higher Education as a Filter," *Journal of Public Economics*, **2**(July 1973).

Barry, M. J., P. D. Cleary, and H. V. Fineberg, "Screening for HIV Infection: Risks, Benefits, and the Burden of Proof," *Journal of Law, Medicine & Health Care*, **14**(December 1986).

Borenstein, S., "Price Discrimination in Free-Entry Markets," *Rand Journal of Economics*, **16**(Autumn 1985).

Brodrick, Mike, and Robert W. Beal, "Blood Testing for Disability Income: Factoring in the AIDS Threat," *Reinsurance Reporter* (Winter 1988).

Chamberlin, J. R., "Assessing the Fairness of Insurance Classifications," *Research in Law and Economics,* **7**(1985).

Clifford, K. A., and R. P. Iuculano, "AIDS and Insurance: The Rationale for AIDS-Related Testing," *Harvard Law Review,* **100**(May 1987).

Crocker, K. J., and A. Snow, "The Efficiency Effects of Categorical Discrimination in the Insurance Industry," *Journal of Political Economy,* **94**(April 1986).

Hoy, M., "Categorizing Risks in the Insurance Industry," *Quarterly Journal of Economics,* **96**(May 1982).

Miike, Lawrence, Office of Technology Assessment, U.S. Congress, "AIDS Antibody Testing," Testimony Before the House Committee on Small Business Subcommittee on Regulation and Business Opportunities.

Mast, Jess L., "HIV Infection: Its Impact on Mortality and Underwriting," *Reinsurance Reporter,* (July 1987).

Meyer, Klemens B., and Stephen B. Pauker, "Screening for HIV: Can We Afford the False Positives?" *New England Journal of Medicine,* **317**(July 23, 1987).

Rothschild, M., and J. Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics,* **90**(November 1976).

Salop, S., "The Noisy Monopolist: Imperfect Information, Price Dispersion and Price Discrimination," *Review of Economic Studies,* **44**(October 1977).

Salop, S., and J. Salop, "Self Selection and Turnover in the Labor Market," *Quarterly Journal of Economics,* **88**(May 1974).

Schatz, B., "The AIDS Insurance Crisis: Underwriting or Overreaching?" *Harvard Law Review,* **100**(May 1987).

Schmalensee, R., "Imperfect Information and the Equitability of Competitive Prices," *Quarterly Journal of Economics,* **98**(August 1984).

Spence, A. M., "Job Market Signalling," *Quarterly Journal of Economics,* **87**(August 1973).

Stiglitz, J. E., "The Theory of 'Screening', Education and the Distribution of Income," *American Economic Review,* **65**(June 1975).

Varian, H., "A Model of Sales," *American Economic Review,* **70**(September 1980).

Weiss, A., "A Sorting-*cum*-Learning Model of Education," *Journal of Political Economy,* **91**(June 1983).

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

Wolpin, K., "Education and Screening," *American Economic Review,* **67**(December 1977).