

The Theory of Risk Classification

by

Keith J. Crocker
Smeal College of Business
The Pennsylvania State University
University Park, PA 16802-3603

and

Arthur Snow
Department of Economics
University of Georgia
Athens, GA 30602-6254

JEL: D82, G22

Keywords: risk categorization, classification, informational asymmetry, information, insurance.

8.1 INTRODUCTION

The efficiency and equity effects of risk classification in insurance markets have been a source of substantial debate, both amongst economists and in the public policy arena.¹ The primary concerns have been the adverse equity consequences for individuals who are categorized unfavorably, and the extent to which risk classification enhances efficiency in insurance contracting. While equity effects are endemic to any classification scheme that results in heterogeneous consumers being charged actuarially fair premiums, whether such classification enhances market efficiency depends on specific characteristics of the informational environment.

In this contribution we set out the theory of risk classification in insurance markets and explore its implications for efficiency and equity in insurance contracting. Our primary concern is with economic efficiency and the role of risk classification in mitigating the adverse selection that arises when insurance applicants are better informed about their riskiness than insurers. We are also interested in the role of classification risk, that is, uncertainty about the outcome of a classification procedure. This uncertainty imposes a cost on risk averse consumers and is thus a potential cause of divergence between the private and social value of information gathering. In addition, the adverse equity consequences of risk classification bear directly on economic efficiency as they contribute to the social cost of classification risk.

8.2 RISK CLASSIFICATION IN THE ABSENCE OF HIDDEN KNOWLEDGE

We begin by considering as a benchmark the case in which both insurers and insurance applicants are symmetrically uninformed about the applicants' propensities for suffering an insurable loss.

8.2.1 Homogeneous Agents

Formally, the insurance environment consists of a continuum of risk averse consumers, each of whom possesses an initial wealth \bar{W} and may suffer a (publicly-observed) loss D with known probability \bar{p} . Each consumer's preferences are represented by the von Neumann-Morgenstern utility function $U(W)$, which is assumed to be strictly increasing and strictly concave, reflecting risk aversion.

A consumer may purchase insurance against the loss by entering into a contract $C \equiv (m, I)$, which specifies the premium m paid to the insurer and the indemnification I received by the insured when the loss occurs. A consumer's expected utility under the insurance contract C is given by

$$V(\bar{p}, C) \equiv \bar{p}U(W_D) + (1 - \bar{p})U(W_N), \quad (1)$$

where $W_D \equiv \bar{W} - m - D + I$ and $W_N \equiv \bar{W} - m$ denote the consumer's state-contingent wealth levels. The expected profit of providing the insurance contract C is given by

$$\pi(\bar{p}, C) \equiv m - \bar{p}I. \quad (2)$$

In order to be feasible, a contract must satisfy the resource constraint

$$\pi(\bar{p}, C) \geq 0, \quad (3)$$

which requires that the premium be sufficient to cover the expected insurance indemnity.

In this setting, an optimal insurance contract is a solution to the problem of maximizing (1) subject to the feasibility constraint (3), which results in full coverage for losses ($I = D$) at the actuarially fair premium ($m = \bar{p}D$). This contract, which is depicted as F in Figure 1, is also the

competitive equilibrium for an insurance market with free entry and exit when all consumers have the same (publicly observed) probability \bar{p} of suffering loss.

8.2.2 Classification with Heterogeneous Agents

We now turn to the case in which both insurers and insurance applicants have access to a costless and public signal that dichotomizes applicants into two groups. After the signal has been observed, a proportion λ of the agents are known to be *high risk* with probability p^H of suffering the loss, while $1-\lambda$ are *low risk* with loss propensity p^L , where $p^H > p^L$ and $\bar{p} = \lambda p^H + (1-\lambda)p^L$. When each individual's type (p^H or p^L) is publicly observable, insurers in a competitive market equilibrium offer full coverage ($I = D$) to all consumers, and charge the actuarially fair premium $m^r = p^r D$ appropriate for the p^r -types. These contracts are depicted as H^* (L^*) for p^H -types (p^L -types) in Figure 1.

Notice that competitive pressures force firms to implement risk classification based upon the insureds' publicly observed characteristic, p^r . Any insurer attempting to offer a contract that would pool both high and low risks (such as F) recognizes that a competitor could offer a profitable contractual alternative that would attract only the low risks. The exodus of low risks caused by such cream-skimming would render the pooling contract unprofitable.

The introduction of symmetric information about risk type accompanied by categorization based on this information increases the utility of some of the insured agents (low risks, who receive L^*), but reduces the utility of others (high risks, who receive H^*) relative to the pre-classification alternative (when both types receive F). From an efficiency perspective, however, the relevant question is whether the insureds *expect* to be better off when moving from a status-quo without information and risk-based categorization to a regime with information and

risk classification. If an individual who is classified as a p^r – type receives the contract C^r , then the expected utility of the insured in the classification regime is

$$E\{V\} \equiv \lambda V^H + (1-\lambda)V^L \quad (4)$$

where $V^i \equiv V(p^i, C^i)$ for $i \in \{H, L\}$. The corresponding resource constraint is

$$\lambda \pi(p^H, C^H) + (1-\lambda) \pi(p^L, C^L) \geq 0, \quad (5)$$

requiring that premiums collected cover expected indemnity payments per capita.

An *efficient classification contract* is a solution to the problem of maximizing (4) subject to (5), which turns out to be the pooling contract, depicted as F in Figure 1, and which provides full coverage at the pooled actuarially fair premium $\bar{p}D$. The intuition behind this result is straightforward. From an ex ante perspective, there are four possible payoff states: The two loss states and the two risk types. Since individuals are risk averse, ex ante expected utility maximization (3) subject to the resource constraint (4) requires equal consumption in all states, and F is the only zero-profit contract with this property.

The technical rationale for this result can be illustrated with reference to Figure 2, which illustrates the utilities possibilities frontier for the classification regime as locus XFY. The concavity of XFY is dictated by the risk aversion of consumers, and movement along the frontier from X towards Y makes L-type (H-types) better (worse) off. From equation (4), we infer that the slope of an indifference curve for the expected utility of an insured confronting classification risk, dV^H/dV^L , is $-(1-\lambda)/\lambda$. By the concavity of U and Jensen's inequality, the pool F is the unique optimum for the consumer anticipating risk classification.

We conclude that the efficient contract in the classification regime ignores the publicly observed signal, and treats all insureds the same independently of their types. Put differently, when information is symmetric between insurers and insureds, uniformed insureds prefer to

remain uninformed if they anticipate that the information revealed will be used to classify the risks. The reason is that the pooling contract F provides full coverage against two types of risk, the *financial risk* associated with the occurrence of the loss state, and the *classification risk* faced by insurance applicants, who may find out that they are high risk. The competitive equilibrium contracts H^* and L^* satisfy the resource constraint (5) and, therefore, are candidate solutions for optimal classification contracts. However, while they provide complete protection from financial risk, they leave consumers wholly exposed to classification risk. Thus, insurers would use public information to classify insurance applicants, even though risk classification based on new information actually reduces efficiency in this setting, and is therefore undesirable.

8.3 RISK CLASSIFICATION IN THE PRESENCE OF HIDDEN KNOWLEDGE

We now turn to an environment in which the individuals to be insured all initially possess private information about their propensities for suffering loss, as in the model introduced by Rothschild and Stiglitz (1976). Each consumer has prior hidden knowledge of risk type, p^H or p^L , but insurers know only that they face a population of consumers in which a proportion λ ($1-\lambda$) have the loss probability p^H (p^L). Given the nature of the informational asymmetry, in order to be attainable a pair of insurance contracts (C^H, C^L) must satisfy the incentive compatibility (self-selection) constraints

$$V(p^\tau, C^\tau) \geq V(p^\tau, C^{\tau'}) \quad \text{for every } \tau, \tau' \in \{H, L\} \quad (6)$$

as a consequence of the Revelation Principle exposted by Myerson (1979) and Harris and Townsend (1981).

In this informationally constrained setting, an efficient insurance contract can be characterized as a solution to the problem of maximizing the expected utility of low-risk consumers $V(p^L, C^L)$ subject to the resource constraint (5), the incentive constraint (6), and a utility constraint on the welfare of high-risk types

$$V(p^H, C^H) \geq \bar{V}^H. \quad (7)$$

As discussed by Crocker and Snow (1985a), a solution to this problem yields full (partial) coverage for H-types (L-types); both the resource constraint (5) and the utility constraint (7) hold with equality; and the incentive constraint (6) binds (is slack) for high (low) risks.

One element of the class of efficient contracts is depicted in Figure 3 as $\{\hat{C}^H, \hat{C}^L\}$. By construction, the locus FA depicts the set of contracts awarded to low risks that, when coupled with a full-insurance contract to which high risks are indifferent, satisfies the resource constraint with equality.² The full class of efficient contracts is obtained by varying the utility level \bar{V}^H in constraint (7). Setting $\bar{V}^H = V(p^H, F)$ yields the first-best pooling allocation F as a solution to the efficiency problem. Setting lower values for \bar{V}^H results in a redistribution away from H-types toward L-types and a solution in which the types receive distinct contracts, as described above, which entail a deductible for L-types and so are strictly second-best. The particular solution depicted in Figure 3, $\{\hat{C}^H, \hat{C}^L\}$, is obtained when the constraint level of utility for the H-types, \bar{V}^H , is set equal to $V(p^H, \hat{C}^H)$ and results in the efficient contract most preferred by the L-type individuals. The allocation $\{\hat{C}^H, \hat{C}^L\}$ will be referred to in the discussion below as the M-W-S allocation.³

Also depicted in Figure 3 is the Rothschild-Stiglitz separating allocation (H^*, A) , which is the Pareto dominant member of the family of contracts that satisfy the incentive constraints (6)

and the requirement that each type of contract break even individually. The Rothschild-Stiglitz allocation is not an element of the (second-best) efficient set when the proportion of H-types (λ) is sufficiently small. Such a situation is depicted in Figure 3, since both types of customers can be made strictly better off at $\{\hat{C}^H, \hat{C}^L\}$ than they would be at $\{H^*, A\}$. In this particular case, all of the efficient contracts involve a cross-subsidy from L-types to H-types. Only when λ is sufficiently large, so that $\{H^*, A\}$ is contained in the class of efficient allocation, is there an efficient contract that does not entail a cross-subsidy. The utility possibilities frontier associated with the solutions to the efficiency problem is depicted in Figure 4. At one end is the utilities distribution associated with the first-best pooling contract F which involves a large cross-subsidy but no inefficiency since the L-types are not subject to a deductible. As one moves along the efficiency frontier toward the point associated with the M-W-S allocation, the degree of cross subsidy is reduced and the amount of inefficiency increases as the L-types are choosing contracts with higher deductibles.⁴

At this juncture, it is useful to elaborate on the differences between the efficiency approach that we have adopted in this chapter, and the equilibrium analyses that have characterized much of the insurance literature. The potential for the non-existence of a Nash equilibrium in pure strategies that was first observed by Rothschild and Stiglitz is an artifact of the incentives faced by uninformed insurers who compete in the offering of screening contracts to attract customers. This result has spawned a substantial body of work attempting to resolve the nonexistence issue, either through the application of non-Nash equilibrium concepts (Wilson (1977); Riley (1979); Miyazaki (1977)) or by considering alternative extensive form models of the insurance process with Nash refinements (Hellwig (1987); Cho and Kreps (1987)).

Unfortunately, the insurance contracts supported as equilibrium allocations generally differ, and depend on the particular concept or extensive form being considered.

In contrast, the characterization of second-best efficient allocations that respect the informational asymmetries of the market participants is straightforward. The model is that of a social planner guided by the Pareto criterion, and who has the power to assign insurance allocations to the market participants.⁵ While the planner is omnipotent, in the sense of having the ability to assign any allocation that does not violate the economy's resource constraints, it is not omniscient, and so is constrained to have no better information than the market participants.⁶ Hence, the issue of how firms compete in the offering of insurance contracts does not arise, since the social planner assigns allocations by dictatorial fiat subject to the (immutable) informational and resource constraints of the economy. This exercise permits an identification of the best outcomes that could, in principle, be attained in an economy. Whether any particular set of equilibrium mechanics can do as well is, of course, a different issue, and one that we consider in more detail in Section 8.5 below.

Finally, as we close this section, notice that risk classification, accomplished through self-selection based on hidden knowledge of riskiness, is required for efficient contracting in this environment. Specifically, with the exception of the first-best pooling allocation F , all efficient allocations are second best, as they entail costly signaling by low-risk types. These consumers retain some risk by choosing a contract that incorporates a positive deductible, but in so doing they are best able to exploit opportunities for risk sharing given the potential adverse selection of low-risk contracts by high-risk consumers.

8.3.1 Categorization Based on Immutable Characteristics

We suppose for the purposes of this section that consumers differ by an observable trait that is immutable, costless to observe, and correlated with (and, hence informative about) the unobservable risk of loss. Examples of such categorizing tools are provided by, but not restricted to, an insured's gender, age or race, which may be imperfectly correlated with the individual's underlying probability of suffering a loss. The interesting question is whether the information available through categorical discrimination, which can be used by insurers to tailor the contracts that are assigned to insureds based upon their observable characteristics, enhances the possibilities for efficiency.

In the first attempt to examine the implications of permitting insurers to classify risks in this environment, Hoy (1982) considered the effects of categorization on market equilibria. Since there was, and still is, little consensus on the identity of the allocations supported by equilibrium behavior, Hoy considered the pure strategy Nash equilibrium of Rothschild and Stiglitz, the "anticipatory" equilibrium of Wilson (1977), and the equilibrium suggested by Miyazaki (1977) which assumes anticipatory behavior but permits cross-subsidization within an insurer's portfolio of contractual offerings. Hoy found that the efficiency consequences of permitting risk classification were ambiguous, depending on the particular equilibrium configuration posited. The primary reason for this ambiguity is that, with the exception of the Miyazaki equilibrium, none of the allocations supported by the equilibrium behaviors considered is guaranteed to be on the efficiency frontier.⁷ Thus, a comparison of the equilibrium allocations pre- and post-categorization provides no insights regarding whether permitting categorization enhances the efficiency possibilities for insurance contracting.

A more fruitful approach is explored by Crocker and Snow (1986), who compare the utilities possibilities frontier for the regime where categorization is permitted to the one in which

it is not. Throughout the remainder of this section, we assume that each insurance applicant belongs either to group A or to group B , and that the proportion of low-risk applicants is higher in group A than in group B . Letting λ_k denote the proportion of H-types in group k , we have $0 < \lambda_A < \lambda_B < 1$, so that group membership is (imperfectly) informative. Assuming that a proportion ω of the population belongs to group A , it follows that $\omega\lambda_A + (1 - \omega)\lambda_B = \lambda$.

Let $C_k \equiv (C_k^H, C_k^L)$ denote the insurance contracts offered to the members of group k . Since insurers can observe group membership but not risk type, the contractual offerings must satisfy separate incentive constraints for each group, that is,

$$V(p^\tau, C_k^\tau) \geq V(p^\tau, C_k^{\tau'}) \text{ for all } \tau, \tau' \in \{H, L\} \quad (8)$$

for each group $k \in \{A, B\}$. In addition, contracts must satisfy the resource constraint

$$\omega[\lambda_A \pi(p^H, C_A^H) + (1 - \lambda_A) \pi(p^L, C_A^L)] + (1 - \omega)[\lambda_B \pi(p^H, C_B^H) + (1 - \lambda_B) \pi(p^L, C_B^L)] \geq 0, \quad (9)$$

which requires that the contracts make zero profit on average over the two groups combined.

To demonstrate that risk categorization may permit Pareto improvements⁸ over the no-categorization regime, it proves useful to consider the efficiency problem of maximizing $V(p^L, C_B^L)$ subject to the incentive constraints (8), the resource constraint (9), and the utility constraints

$$V(p^\tau, C_A^\tau) \geq V(p^\tau, \hat{C}^\tau) \text{ for } \tau \in \{H, L\}; \text{ and} \quad (10)$$

$$V(p^H, C_B^H) \geq V(p^H, \hat{C}^H), \quad (11)$$

where $\hat{C} \equiv (\hat{C}^H, \hat{C}^L)$ is an efficient allocation in the no-categorization regime. By construction, we know that this problem has at least one feasible alternative, namely the no-categorization contract \hat{C} which treats the insureds the same independently of the group (A or B) to which they

belong. If \hat{C} is the solution, then the utilities possibilities frontier for the categorization and the no-categorization regimes coincide at \hat{C} . However, if \hat{C} does not solve the problem, then categorization admits contractual opportunities Pareto superior to \hat{C} and the utilities possibilities frontier for the categorization regime lies outside the frontier associated with the no-categorization regime.

Let δ denote the Lagrange multiplier associated with the utility constraint (7) for the efficiency problem in the no-categorization regime, and let μ_H be the multiplier associated with the incentive constraint (6) for $\tau = H$. The following result is from Crocker and Snow (1986, p. 329).

Result: Categorization permits a Pareto improvement to be realized over efficient contracts without categorization if and only if

$$\frac{\delta}{\mu_H} < \frac{\lambda - \lambda_A}{\lambda_A(1 - \lambda)}. \quad (12)$$

For the inequality to hold, it is sufficient that $\delta = 0$, which necessarily obtains whenever the utility constraint, \bar{V}^H , in (7) is set sufficiently low. When $\delta > 0$, the location of the utilities possibilities frontiers depends on the informativeness of the categorization. When categorization is more informative, λ_A is smaller and the right hand side of (12) is larger. If categorization were uninformative ($\lambda = \lambda_A$), then (12) could never hold, and if categorization were perfectly informative ($\lambda_A = 0$), then (12) would always be satisfied. Finally the inequality can never hold when $\mu_H = 0$, which occurs when the incentive constraint (6) for the efficiency problem in the no-categorization regime is slack. Contract F is the only efficient contract for which the incentive constraint is slack, so that the utilities possibilities frontiers always coincide at F

regardless of the degree of informativeness of the categorization. The relative positions of the utilities possibilities frontiers for the categorization and the no-categorization regimes for those in group A are depicted in Figure 5, while a similar diagram applies to those in group B .

To evaluate the efficiency of categorization, we employ the Samuelson (1950) criterion for potential Pareto improvement. Risk classification through *a priori* categorization by insurers is defined to be efficient (inefficient) if there exists (does not exist) a utility distribution in the frontier for the no-categorization regime Pareto dominated by a distribution in the frontier for the categorization regime, and there does not exist (exists) a distribution in the categorization frontier Pareto dominated by one in the no-categorization frontier. Since costless categorization shifts outward the utilities possibilities frontier over some regions and never causes the frontier to shift inward, we conclude that categorization is efficient.

Crocker and Snow (1985b) show that omniscience is not required to implement the hypothetical lump-sum transfers needed to effect movement along a utilities possibilities frontier. Although the appropriate lump-sum transfers cannot be assigned directly to individual consumers, since their risk types are hidden knowledge, these transfers can be built into the premium-indemnity schedule so that insurance applicants self-select the taxes or transfers intended for their individual risk types. In this manner, a government constrained by the same informational asymmetry confronting insurers can levy taxes and subsidies on insurance contracts to implement redistribution, while obeying incentive compatibility constraints and maintaining a balanced public budget. Our application of the Samuelson criterion is thus consistent with the informational environment.

The situation is somewhat different when consumers differ by an observable, immutable trait that is correlated with the unobservable risk of loss, but is costly to observe. Crocker and

Snow (1986) show that the utilities possibilities frontiers cross in this case, so long as the cost is not too high. Intuitively, the cost of categorization amounts to a state-independent tax on each consumer's wealth. As a result, when the adverse selection externality is not very costly and low-risk types are nearly fully insured, categorization costs dominate the small efficiency gains realized by the winners leaving no possibility of compensating the losers. Conversely, if the adverse selection externality imposes sufficient cost on the low-risk consumers, then gains from categorization realized by the winners are sufficient for potential Pareto improvement provided categorization is not too costly. This situation is depicted in Figure 6.

If categorization were required, then insurers would sometimes categorize insurance applicants even when the result is not a potential Pareto improvement over not categorizing. In this scenario the efficiency effects of costly categorization are ambiguous. As Rothschild (2011) points out, however, the second-best efficient allocations when categorizing is costly do not require the use of categorization. Consider a social planner with the power to assign insurance contracts to applicants subject to the economy's resource and informational constraints, and who has access to the same costly categorizing technology as insurers. Because the social planner can choose not actually to employ the categorizing technology, the second-best Pareto frontier for the planner is the outer envelope of utility possibilities. The Samuelson criterion therefore leads to the conclusion that allowing costly categorization is more efficient than either banning or requiring categorization.

Rothschild further shows that this application of the Samuelson criterion is again consistent with the informational environment. Specifically, for any allocation in the no-categorization regime, a government constrained by the same informational asymmetry confronting insurers can simultaneously provide a properly calibrated social insurance policy and

also legalize categorization so that, in response, insurers choose to employ categorization precisely when doing so yields a Pareto improvement over not categorizing. In this sense, the no-categorization regime is inefficient.

8.3.2 An Empirical Estimate: The Case of Annuities

Finkelstein, Poterba and Rothschild (2009) adapt the basic framework of Hoy (1982) and Crocker and Snow (1986) to facilitate empirical estimates of the efficiency and distributional consequences of prohibiting categorical discrimination in real-world insurance markets. Their approach is to use an empirically calibrated model to estimate the welfare consequences restricting gender-based pricing in the compulsory annuities market of the United Kingdom. In this market, which is described in greater detail in Finkelstein and Poterba (2002, 2004), retirees are required to annuitize a substantial portion of their accumulated tax-preferred retirement savings, but there is scope for annuity providers to screen different risk types by offering annuity contracts with different lifetime payout structures.

Their adaptation requires two significant modifications of the standard insurance model. First, the model is extended to allow many “indemnity” states that correspond to annuity payments in future years, where the uncertainty arises because the annuity is paid only if the annuitant survives. From the insurer’s perspective, low-risk (high-risk) individuals are those that have a lower longevity (higher longevity), and this is assumed to be private information known only to the annuitant. Second, the model allows for the possibility that individuals could, in a fashion that is not observable to the insurer, save a portion of their annuity income to supplement the consumption provided by the annuity at later ages, in effect, permitting individuals to engage in a form of “self-insurance”.

Using mortality data from a major insurer, maximum likelihood estimation is used to calibrate a model with two unobservable types (high-risk and low-risk) and two observable categories (male and female). The categories are observable to the insurer and each category contains both high- and low-risk types, although the female category contains a higher proportion of high-risk (longer-lived) individuals. When insurers are permitted to categorize their insurance offerings on observable gender, the market segments into male and female sub-markets in which insurers screen each category for unobservable risk type through their contractual offerings. The result is screening of types in both gender categories in the manner of Figure 3, but with different contracts offered to male and female applicants. In contrast, when such categorical discrimination is prohibited, insurers still screen types as in Figure 3, but now are constrained to offer the same screening contracts to both genders. As a result, when calculating the efficiency costs of prohibiting gender-based pricing, there are in principle three efficiency frontiers that must be considered: those associated with each of the two genders when categorical discrimination is permitted, and the one associated with the regime in which such discrimination is prohibited.

The goal in Finkelstein *et al.* is to calculate bounds on the welfare costs associated with a ban on gender-based pricing. Their approach is to assume that, when gender discrimination is allowed, the segmented markets provides a second-best efficient allocation to each category, and that there is no cross-subsidy between the two observable categories. In contrast, when gender-based pricing is banned, the market is assumed to attain an allocation on a no-categorization efficiency frontier of the type described by Figure 4. As noted by Crocker and Snow (1986, p. 329), starting from an efficient contract on the no-categorization frontier, it is possible to make the category composed of fewer high risks better off, and at a lower resource cost, if risk categorization were to be introduced. This saving in resources represents the efficiency cost of the categorization ban. Thus, the potential efficiency cost of

a ban on gender-based pricing ranges from zero if the post-ban market achieves the first best pooling allocation F (which results in maximal across-gender redistribution) to its maximum value when the post-ban result is the M-W-S allocation (which result in the minimal across-gender redistribution).

Figure 7 (which is Figure 4 from Finkelstein *et al.*) depicts the efficient annuity contracts associated with the W-M-S allocation in the presence of a ban on gender-based pricing. High-risk (long-lived) types receive a full insurance annuity that provides constant real payments for the duration of their retirements. Low-risk types, by contrast, receive a front-loaded annuity. This front loading allows them to receive substantially higher annuity payments for most of their expected lifetimes while still effectively discouraging the high-risks from selecting the annuity targeted to the low-risk types. Moreover, the efficient annuities involve a cross-subsidy from low- to high-risk types since the latter obtain a better than actuarially fair annuity payment. Since the high-risks are the recipients of the subsidy, and the female category contains a disproportionate share of the high-risk annuitants, the effect is to generate a cross subsidy from males to females. Column (9) of the Table (which is Table 3 from Finkelstein *et al.*) quantifies the cross-subsidy associated with the post-ban W-M-S allocation, which is on the order of a two to four per-cent transfer of the retirement wealth, depending on the degree of risk aversion. As one moves along the utility possibilities frontier, the size of this cross-subsidy increases and achieves its maximum at the full insurance pooling allocation F , which is reported in column (10) as 7.14%.

The Table also quantifies the efficiency costs associated with a ban on the gender-based pricing of annuities. The maximum efficiency cost occurs if the post-ban market achieves the M-W-S allocation, which is column (5) of the Table and results in an efficiency cost ranging from .018-.025% of retirement wealth, depending on the degree of risk aversion. Other post-ban allocations result in lower efficiency costs, and the first-best pooling contract (point F on the efficiency frontier) results in

no efficiency cost, as reported by column (6). While the efficiency costs of the ban on gender-based pricing are non-zero, as predicted by Crocker and Snow (1986), they are small relative to the degree of redistribution effected by the ban.

8.3.3 Categorization Based on Consumption Choices

In contrast to categorical discrimination based on observable but immutable characteristics, in many situations consumers use products, such as cigarettes or stodgy automobiles, with the anticipation that such consumption will affect their opportunities for insuring. The actuarial relationship between the consumption of such a *correlative product* and underlying risk may be the consequence of a direct causal link (smoking and heart disease) or merely a statistical relationship (people who drive stodgy automobiles are more likely to be careful drivers). In both cases, however, the observed consumption of a correlative product permits insurers to design contracts that mitigate the problems of moral hazard and adverse selection inherent in insurance markets with private information.

To analyze the efficiency effects of permitting insurers to classify applicants on the basis of their consumption choices, Bond and Crocker (1991) assume that consumers' utility functions have the additively separable form

$$U(W) + \theta G(x) \tag{13}$$

where W and x are the consumer's wealth and consumption of the correlative product, respectively, and θ is a taste parameter. There are two types of consumers distinguished by their

taste for the correlative product $\theta \in \{\theta^H, \theta^L\}$ where $\theta^H > \theta^L$. The proportion of θ^H -types in the population is λ .

Each consumer faces two possible wealth states, so W_D (W_N) represents consumption of other goods (that is, wealth net of expenditures on the correlative productive) in the loss (no-loss) state. The probability of the loss state for a θ^r -type consumer is $p^r(x)$, with $\partial p^r(x)/\partial x \geq 0$ and $1 \geq p^H(x) \geq p^L(x) \geq 0$ for every x . Thus, the consumption of the correlative product either affects directly, or may be positively correlated with, the potential for loss. While we restrict our attention to the case of hazardous goods whose level of consumption increases the probability of a loss ($\partial p^r/\partial x > 0$) or where the consumer's taste for the product is positively correlated with loss propensity ($p^H(x) > p^L(x)$), consideration of other correlative relationships is straightforward.

Under the assumption that consumers purchase the hazardous good x before the wealth state is revealed, the expected utility of a type θ^r individual is

$$V^r(W_D, W_N, x) \equiv p^r(x)U(W_D) + (1-p^r(x))U(W_N) + \theta^r G(x). \quad (14)$$

When the hazardous good is supplied by a competitive market at marginal cost c , the state-contingent wealth of an insured is now $W_N \equiv \bar{W} - m - cx$ and $W_D \equiv \bar{W} - m - cx + I - D$. The expected profit of providing the insurance policy $\{m, I\}$ to a θ^r -type agent who consumes x is

$$\pi^r(m, I, x) \equiv m - p^r(x)I. \quad (15)$$

A contract $C \equiv \{m, I, x\}$ determines the consumption bundle for the insured, and an *allocation* (C^H, C^L) is a pair of contracts assigned to insureds based upon their types. Feasible contracts must satisfy the resource constraint

$$\lambda \pi^H(C^H) + (1-\lambda) \pi^L(C^L) \geq 0, \quad (16)$$

which ensures that premiums are sufficient to cover expected indemnity payments per capita.

When the insureds' taste parameters and the consumption of the hazardous good can be observed publicly, first-best allocations are attainable. In that event, an efficient allocation, denoted (C^{L*}, C^{H*}) , is a solution to the problem of maximizing $V^L(C^L)$ subject to (16) and a utility constraint on H-types, $V^H(C^H) \geq \bar{V}^H$. An efficient allocation results in full insurance ($W_D^\tau = W_N^\tau = W^\tau$) for both types of agents, and consumption levels for the hazardous good, x^τ , that equate each type of consumer's marginal valuation of consumption with its marginal cost, that is,

$$\frac{\theta^\tau G'(x^\tau)}{U'(W^\tau)} = c + D\partial p^\tau(x^\tau) / \partial x, \quad (17)$$

Notice that the marginal cost of the hazardous good includes its production cost c as well as its marginal effect on the expected loss.

The interesting case from the perspective of risk classification arises when consumption of the hazardous good, x , is observable but the consumer's taste parameter, θ , is private information. In this setting with asymmetric information, allocations must satisfy the incentive constraints

$$V^\tau(C^\tau) \geq V^\tau(C^{\tau'}) \text{ for all } \tau, \tau' \in \{H, L\}. \quad (18)$$

This case is referred to as *endogenous risk classification* since the consumers' insurance opportunities may depend on their choices regarding consumption of the hazardous good.

An efficient allocation is a solution to the problem of maximizing $V^L(C^L)$ subject to $V^H(C^H) \geq \bar{V}^H$, the incentive constraints (18), and the resource constraint (16). There are two classes of solutions, which differ based on whether any of the incentive constraints (18) are binding.

8.3.4 First-Best Allocations: A Pure Strategy Nash Equilibrium

When the incentive constraints (18) do not bind at a solution to the efficiency problem, the efficient allocation provides full coverage to all individuals and charges actuarially fair premiums $p^r(x^r)D$ that depend on the amount of the hazardous good consumed (as determined by (17)). The insurance premium offered is bundled with a consumer's observed consumption of the hazardous good, so that individuals are classified based upon their consumption choices for x . An efficient allocation in this case is depicted as (C^H*, C^L*) in Figure 8.

The moral hazard aspect of hazardous goods consumption is reflected by the curvature of a consumer's budget constraint $W = \bar{W} - p^r(x)D - cx$, which reflects the fact that the risk of loss depends on consumption of the hazardous good, given $\partial p^r(x)/\partial x \neq 0$. The potential for adverse selection arises because the budget constraint for θ^H -types lies below that for θ^L -types, since $p^H(x) > p^L(x)$. In the special case where there is no adverse selection ($p^H(x) = p^L(x)$), the budget constraints of the two types of consumers coincide, and a first-best allocation solves the efficiency problem. Effectively, the insurer levies a Pigovian tax based upon the observed consumption levels of the hazardous good, thereby forcing the insured to internalize the moral hazard externality. Introducing a small amount of private information still permits the attainment of first-best allocations, as long as the difference in loss probabilities ($p^H(x) - p^L(x)$) is not too great.

It is easy to see that the first-best allocation (C^H*, C^L*) is necessarily a Nash equilibrium in pure strategies whenever the incentive constraints (18) are not binding. This result provides an important insight concerning the desirability of permitting insurers to classify applicants on the basis of their consumption of goods that directly affect loss propensities. In the polar case, where the level of hazardous good consumption completely determines an individual's loss

probability (so $p^H(x) = p^L(x) \equiv p(x)$), endogenous risk classification allows first-best allocations to be attained as Nash equilibria. Indeed, to disallow such categorization would cause a reversion to the typical adverse selection economy where the Nash equilibrium, if it exists, lies strictly inside the first-best frontier.

Even in cases where endogenous risk classification is imperfect, so that some residual uncertainty about the probability of loss remains after accounting for consumption of the hazardous good ($p^H(x) \neq p^L(x)$), the pure strategy Nash equilibrium exists and is first-best efficient as long as the risk component unexplained by x is sufficiently small. Consequently, insurers may alleviate the problems of adverse selection in practice by extensively categorizing their customers on the basis of factors causing losses, which may partly offset the insureds' informational advantage and permit the attainment of first-best allocations as equilibria.

8.3.5 Second-Best Allocations

When incentive constraints are binding at a solution to the efficiency problem, an optimal allocation generally results in distortions in both the insurance dimension and in the consumption of the hazardous good. While the nature of a second-best allocation depends on the specifics of the model's parameters, there are several generic results.

Result. When the incentive constraint (18) binds for the θ^H -type consumer, an efficient allocation is second best. Also,

- (i) if $p^H(x) > p^L(x)$, then θ^H -types (θ^L -types) receive full coverage (are under-insured); and

$$(ii) \quad \text{if } \left\{ \begin{array}{l} \text{either } p^H(x) = p^L(x) \text{ (no adverse selection case)} \\ \text{or } \frac{\partial p^r(x)}{\partial x} = 0 \text{ (pure adverse selection case) and } \frac{\theta^H}{\theta^L} = \frac{p^H}{p^L} \end{array} \right\}$$

r

then θ^L -types (θ^H -types) under-consume (over-consume) the hazardous good relative to the socially optimal level (17).

These results indicate the extent to which there is a tension between discouraging consumption of the hazardous good to mitigate moral hazard, on the one hand, and using such consumption as a signal to mitigate adverse selection, on the other. An optimal contract reflects a balance between the signaling value of hazardous goods consumption, and the direct social costs imposed by the consumption of products that increase the probability of loss.

As an example, consider those who ride motorcycles without wearing safety helmets, which is a form of hazardous good consumption. On the one hand, those who choose to have the wind blowing through their hair are directly increasing their probabilities of injury (the *moral hazard* effect), which increases the cost of riding motorcycles. On the other hand, the taste for not wearing helmets may be correlated with a propensity of the rider to engage in other types of risk-taking activities (the *adverse selection* effect), so that the choice to ride bear-headed may be interpreted by insurers as an imperfect signal of the motorcyclist's underlying risk. Interestingly, to require the use of safety helmets eliminates the ability of insurers to utilize this signal, with deleterious effects on efficiency.

8.4 RISK CLASSIFICATION AND INCENTIVES FOR INFORMATION GATHERING

As discussed originally by Dreze (1960) and subsequently by Hirshleifer (1971), because information resolves uncertainty about which of alternative possible outcomes will occur, information destroys valuable opportunities for risk averse individuals to insure against unfortuitous outcomes. This phenomenon lies behind the observation, made earlier in section

8.2.2, that new information used by insurers to classify insurance applicants has an adverse effect on economic efficiency. As emphasized in the “no-trade” theorem of Milgrom and Stokey (1982), if applicants were able to insure against the possibility of adverse risk classification, then new information would have no social value, either positive or negative, as long as consumers initially possess no hidden knowledge.

By contrast, the results of Crocker and Snow (1986) and Bond and Crocker (1991) show that new information can also create valuable insurance opportunities when consumers are privately informed. Information about each consumer’s hidden knowledge, revealed by statistically correlated traits or behaviors, allows insurers to sort consumers more finely, and thereby to reduce the inefficiency caused by adverse selection. In this section, we investigate the effects of risk classification on incentives for gathering information about underlying loss probabilities.

8.4.1 Symmetric Information

Returning to the benchmark case of symmetric information, we now suppose that some consumers initially possess knowledge of being either high-risk or low-risk, while other consumers are initially uninformed. Being symmetrically informed, insurers can classify each insurance applicant by informational state and can offer customers in each class a contract that provides full coverage at an actuarially fair premium. Thus, with reference to Figure 1, informed consumers receive either H^* or L^* , while uninformed consumers receive the first-best pooling contract F .

Observe that uninformed consumers in this setting have no incentive to become informed, since they would then bear a classification risk. In Figure 2, the line tangent to the utilities possibilities frontier at point F corresponds to an indifference curve for an uninformed

consumer.⁹ Clearly, the pooling contract F is preferred to the possibility of receiving H^* with probability λ or L^* with probability $1 - \lambda$, that is,

$$V(\bar{p}, F) > \lambda V(p^H, H^*) + (1 - \lambda)V(p^L, L^*),$$

where $\bar{p} = \lambda p^H + (1 - \lambda)p^L$. Since all three of the contracts (F, L^*, H^*) fully insure consumers against the financial risk associated with the loss D , becoming informed in this environment serves only to expose a consumer to classification risk, with no countervailing gain in efficiency. The incentive for uninformed consumers to remain uninformed is consistent with socially optimal information gathering, since the classification risk optimally discourages individuals from seeking information.

8.4.2 Initial Acquisition of Hidden Knowledge

Hidden knowledge can be acquired either purposefully or serendipitously as a by-product of consumption or production activities. In this section we consider environments in which some consumers initially possess *hidden knowledge* of their riskiness, while others do not. Moreover, we assume that insurers cannot ascertain *a priori* any consumer's informational state. Figure 9 illustrates the Pareto dominant separating allocation in which each contract breaks even individually, which is the analogue to the Rothschild and Stiglitz equilibrium with three types $(p^H, \bar{p} \text{ and } p^L)$ of consumers.¹⁰ Consumers with hidden knowledge of risk type (either p^H or p^L) select contract H^* or contract L , while those who are uninformed (perceiving their type to be \bar{p}) select contract B on the pooled fair-odds line. Notice that the presence of uninformed consumers adversely affects low-risk types, who could otherwise have received the (preferred) contract A . Thus, the presence of uninformed consumers may exacerbate the adverse selection inefficiency caused by the hidden knowledge of informed consumers.

In this setting, and in contrast to the case of symmetric information in 8.4.1 above, uninformed consumers *do* have an incentive to become informed despite the classification risk they must bear as a result. Ignoring any cost of acquiring information, and assuming for the moment that contracts H^* and L continue to be offered, the expected gain to becoming informed is given by

$$\lambda V(p^H, H^*) + (1 - \lambda)V(p^L, L) - V(\bar{p}, B) = (1 - \lambda)[V(p^L, L) - V(p^L, B)],$$

where the equality follows from the fact that

$$V(\bar{p}, B) \equiv \lambda V(p^H, B) + (1 - \lambda)V(p^L, B),$$

and from the binding self-selection constraint requiring that $V(p^H, H^*) = V(p^H, B)$. The incentive constraints also require that $V(p^L, L)$ exceeds $V(p^L, B)$. Hence, for an uninformed consumer, the expected gain in utility to becoming informed of risk type (p^H or p^L) is unambiguously positive. Finally, when all consumers possess hidden knowledge, contract A replaces contract L , which enhances the expected value of becoming informed, while also raising the utility of low-risk insureds. We conclude that, in the presence of adverse selection, risk classification through self-selection provides an incentive for uninformed consumers to acquire hidden knowledge, and that this action enhances the efficiency of insurance contracting by reducing, in the aggregate, the amount of signaling required to effect the separation of types.

This result strengthens the finding reported by Doherty and Posey (1998), who adopt the additional assumption that high-risk consumers, whose test results have indicated a risk in excess of p^H , can undergo a treatment that reduces the probability of loss to p^H . They emphasize the value of the treatment option in showing that initially uninformed consumers choose to acquire hidden knowledge. Our demonstration of this result abstracts from the possibility of treatment, and reveals that risk classification is valuable to uninformed consumers in markets where some

consumers possess hidden knowledge, despite uncertainty about the class to which one will be associated. Thus, private incentives for information gathering accurately reflect the social value of initially acquiring hidden knowledge.

A case of special concern arises when information reveals whether a loss has occurred, as when an incurable disease is diagnosed. Figure 10 illustrates this situation with $p^H = 1$ and $p^L = 0$. The equilibrium indifference curve for H-type consumers coincides with the forty-five degree line, while that for L-types coincides with the horizontal axis. Although informed consumers possess no insurable risk, uninformed consumers do possess an insurable risk. However, when insurers are unable to distinguish between insurance applicants who are informed and those who are not, the market fails to provide any insurance whatsoever.¹¹ This result, obtained by Doherty and Thistle (1996), represents the extreme case in which uninformed consumers have no incentive to acquire hidden knowledge. Notice that the acquisition of such knowledge has no social value as well, so that private incentives are once again in accord with economic efficiency.

8.4.3 Acquisition of Additional Hidden Knowledge

Henceforth, we assume that all consumers possess hidden knowledge. In this section, we investigate the private and social value of acquiring additional hidden knowledge. Since hidden knowledge introduces inefficiency by causing adverse selection, it is not surprising to find that additional hidden knowledge can exacerbate adverse selection inefficiency. However, we also find that additional hidden knowledge can expand opportunities for insuring, and thereby mitigate adverse selection inefficiency.

We assume that all insurance applicants have privately observed the outcome of an experiment (the α -experiment) that provides information about the underlying probability of loss, and we are concerned with whether the acquisition of additional hidden knowledge (the β -

experiment) has social value. Prior to observing the outcome of the α -experiment, all consumers have the same prior beliefs, namely that the loss probability is either p^1 or p^2 ($> p^1$) with associated probabilities denoted by $P(p^1)$ and $P(p^2)$ such that

$$\bar{p} = p^1 P(p^1) + p^2 P(p^2).$$

After the α -experiment, consumers who have observed $\alpha^r \in \{\alpha^L, \alpha^H\}$ have formed posterior beliefs such that

$$p^r = p^1 P(p^1 | \alpha^r) + p^2 P(p^2 | \alpha^r).$$

A proportion $\lambda = P(\alpha^H)$ have observed α^H .

At no cost, consumers are permitted to observe a second experiment (the β -experiment) whose outcome $\beta^i \in \{\beta^1, \beta^2\}$ reveals the consumer's actual loss probability $p^i \in \{p^1, p^2\}$. In what follows, the notation $P(\beta^i, \alpha^j)$ denotes the joint probability of observing the outcome (β^i, α^j) of the two experiments, where $i \in \{1, 2\}$ and $j \in \{H, L\}$, while $P(\beta^i)$ denotes the marginal probability $P(\beta^i, \alpha^L) + P(\beta^i, \alpha^H)$.

For this environment, Crocker and Snow (1992) establish the following propositions concerning the efficiency implications of the additional hidden knowledge represented by the second experiment β . The experiment has a positive (negative) social value if the utilities possibilities frontier applicable when consumers anticipate observing β prior to contracting lies (weakly) outside (inside) the frontier applicable when observing β is not an option.

Result: The additional hidden knowledge represented by experiment β has a positive social value if

$$p^2 P(\beta^2, \alpha^L) - p^1 P(\beta^1, \alpha^H) \leq \min\{P(\beta^2, \alpha^L) - P(\beta^1, \alpha^H), P(\beta^2)(p^2 - p^1)/(1 - p^H)\},$$

but has a negative social value if

$$p^2 P(\beta^2, \alpha^L) - p^1 P(\beta^1, \alpha^H) \geq \max \{0, [p^2 P(\beta^2) - p^H P(\alpha^H)]/p^H\}.$$

So, for example, if the probability difference $P(\beta^2, \alpha^L) - P(\beta^1, \alpha^H)$ is positive, then the weighted difference $p^2 P(\beta^2, \alpha^L) - p^1 P(\beta^2, \alpha^H)$ cannot be too large, for then the acquisition of the hidden knowledge β would have negative social value. Similarly, if the probability difference is negative, then the weighted difference must also be negative in order for β to have positive social value. Although these conditions are not necessary for additional hidden knowledge to have a positive or negative social value, they depend only on exogenous parameters of the informational environment without regard to consumers' risk preferences.

Figure 11 illustrates the sources of social gains and losses from additional hidden knowledge. In the absence of experiment β , a typical efficient separating allocation is depicted by the pair (H^*, A) . Once consumers have privately observed β , the pair (H^*, A) is no longer incentive compatible. The α^L -type consumers who discover their type to be p^2 now prefer H^* to their previous allocation A , while the α^H -types who find out that their loss propensity is p^1 now prefer A . The effect of consumers' acquiring additional hidden knowledge through the β -experiment is to alter irreversibly the set of incentive compatible allocations, and to render previously feasible contracts unattainable. From a social welfare perspective, for the β -experiment to have positive social value, there must exist allocations that (i) are incentive compatible under the new (post β -experiment) informational regime, (ii) allow consumers to be expectationally at least as well off as they were at (H^*, A) prior to the experiment; and (iii) earn nonnegative profit.

It is easy to verify that the incentive compatible pair (\hat{H}, A) , when evaluated by consumers ex ante, prior to observing β , affords α^L -types (α^H -types) the same expected utility

they enjoy at A (H^*).¹² Notice that α^L -types who observe β^2 no longer bear signaling costs since they no longer choose the deductible contract A , while α^H -types who observe β^1 now absorb signaling costs. Since, by construction, consumers are indifferent between not observing the β -experiment and receiving (H^*, A) , or observing the β -experiment and being offered (\hat{H}, A) , the acquisition of the additional hidden information has positive social value if the contracts (\hat{H}, A) yield positive profit to the insurer.¹³ Whether this occurs depends on the proportion of consumers signaling less when newly informed, $p^2P(\beta^2, \alpha^L)$, relative to the proportion signaling more, $p^1P(\beta^1, \alpha^H)$, as indicated by conditions stated in the *Result* above.

Private incentives for information gathering may not accord with its social value in the present environment. We will illustrate this result in a setting where insurance markets attain separating equilibria in which contracts break even individually. First, notice that, if α^L -types acquire access to the β -experiment, then α^H -types prefer also to become informed, even though they may be worse off than if neither type has access to the β -experiment. To see this, refer to Figure 12, which illustrates the equilibrium when only α^L -types will observe β and receive either H^2 or L , and α^H -types will not observe β and bear adverse selection costs by receiving H instead of H^* . The α^H -types would be indifferent between remaining uninformed and receiving H , or observing β and afterwards selecting either H^2 or H , since

$$P(\beta^2|\alpha^H)V(p^2, H^2) + P(\beta^1|\alpha^H)V(p^1, H) = V(p^H, H)$$

given the equality $V(p^2, H^2) = V(p^2, H)$ implied by incentive compatibility. Moreover, it follows that α^H -types would strictly prefer to observe β and afterwards select H^2 or A^1 , even though they may be worse off than they would have been receiving H^* , which is rendered unattainable once

α^L -types have private access to experiment β . Thus, once the α^L -types become informed, it is in the best interests of α^H -types to do so as well.

Second, note that α^L -types will demand the β -experiment even if their gains are negligible and are more than offset by the harm imposed on α^H -types, so that the social value of the β -experiment is negative. To demonstrate this result, refer to Figure 13 which illustrates a “knife-edge” case where α^L -types are just indifferent to acquiring additional hidden knowledge.¹⁴ The α^H -types, however, are necessarily worse off, since

$$V(p^H, H^*) > V(p^H, A^1) = P(\beta^2 | \alpha^H) V(p^2, H^2) + P(\beta^1 | \alpha^H) V(p^1, A^1),$$

where the equality follows from the self-selection condition $V(p^2, H^2) = V(p^2, A^1)$. If α^L -types were to experience a small expected gain from acquiring additional hidden knowledge, they would demand access to the β -experiment even though this information would be detrimental to efficiency in insurance contracting. In such an environment, private incentives for information gathering do not reflect its social value. The problem is that the acquisition of private information by some consumers generates an uncompensated externality for others through its effect on the incentive constraints.

8.4.4 Acquisition of Private Information: The Case of Genetic Tests

There are substantial differences worldwide in how countries regulate (or not) the use of genetic testing for life, private health, and long term disability insurance purposes. Regulation varies from total to partial legislative prohibition banning the use of genetic test results by insurers, on the one hand, to voluntary moratoria or laissez faire with no regulatory or voluntary restrictions, on the other. In most of Western Europe the ban is almost total, falling in line with the UNESCO Declaration on Human Genetic Data 2003. In Belgium, insurers are prohibited

from even accepting favourable genetic test results provided voluntarily by consumers. In the United Kingdom and the Netherlands, companies can ask for genetic test results only for large policies (those exceeding £500,000 in Britain and, for life insurance in the Netherlands, policies exceeding 300,000 Dutch guilders.¹⁵ The latter is adjusted every three years to the cost of living.) In Britain, the types of genetic tests that insurers can request for policies exceeding the cap are restricted to tests deemed relevant by an independent committee. Australia, New Zealand and Canada are among those who have not introduced any legislation. The United States is a particular case in that the discussion there, in the absence of socialized medical insurance, involves both the health and the life insurance industries, and the regulations vary from state to state. Federally, the Genetic Information Nondiscrimination Act (GINA) passed in May of 2008 addresses the use of genetic testing in health insurance although only 14 states have introduced some laws to govern the use of genetic testing in life insurance and these laws generally entail restrictions rather than outright bans.¹⁶

Restricting the use of genetic test results for pricing life insurance or annuity products seems likely to have the potential for strong adverse selection effects. However, much less is known about the actuarial relevance of genetic test results in the population. Most people do not currently possess genetic test results and those that have strong actuarial impact, such as the Huntington Disease gene, are very rare, approximately one in 10,000. Most empirical and simulation analysis to date suggest that restricting insurers' use of genetic test results for ratemaking purposes is unlikely to have a significant impact on insurance markets. In a review of the current actuarial (academic) literature MacDonald (2009, p. 4) concludes that "little, if any strong empirical evidence has been found for the presence of adverse selection (although it is admittedly hard to study)." Simulation exercises based on population genetics and epidemiological data by Hoy, et al.

(2003) and Hoy and Witt (2007) also find, for the most part, little impact is likely to occur from a ban on insurers using genetic test results for health and life insurance markets, respectively. Oster, et al. (2010), however, report “strong evidence of adverse selection” in the long-term care insurance market due to individuals holding private information about their Huntington Disease carrier status. Moreover, Hoy and Witt (2007) demonstrate that the effect of adverse selection on market behavior for many diseases is likely to depend on family history. This follows since those with a family history of the particular disease – even those associated with so-called predisposition genes such as the BRCA1/2 genes – are more likely to obtain a genetic test and are more likely to carry the gene. Those who do possess the relevant gene are even more likely to incur the disease than their family history alone would suggest. Table 2, which is based on two of the 13 different types of family history analyzed in Hoy and Witt (2007) illustrates this point. The “low family risk” background reflects women (aged 35 to 39 years) who have no family background of breast or ovarian cancer (from mother or sisters), while the “high family risk” background reflects women who have had a mother who had ovarian cancer before age 50 as well as breast cancer (after age 50). The unconditional probability of incurring breast cancer within 10 years for a woman with the high-risk family background is greater by a factor of approximately 2.2 and the probability of the high risk woman having one of the BRCA1/2 mutations present is higher by a factor of 65. The table also reports the probability of breast cancer conditional on the result of a genetic test and the probability is much higher for both family backgrounds if one of the mutations is present.

In a simulation model of 10-year term life insurance purchases, which is based on socioeconomic factors as well as various assumptions regarding the degree of risk aversion, it turns out that if 100% of all women in each risk group were to obtain a genetic test for one of the

BRCA1/2 genes and were allowed to keep that information private, then the effect on price would be about a 1.5% increase for the low family risk type but almost a three-fold increase in price for those with the high-risk background. This result demonstrates how sensitive the market reaction may be to the fraction of people who may hold important genetic information. It also raises interesting questions about the use of family history – which is at least in part crude genetic information – as a relevant and allowed categorical variable. Given the potential growth of genetic information among the public through direct-to-consumer testing services, the future holds substantial uncertainty in this regard and continued empirical research will be necessary in order to help resolve the debate about the use of genetic information in insurance markets.

There are two popular and conflicting views of the importance of discrimination in deciding on whether insurers should be allowed to use genetic test results (or family background of diseases for that matter) in pricing insurance. The “actuarial” view is that any actuarially relevant information should be allowed, even if imperfect, and it is not discriminatory to charge people who impose higher expected costs on insurers a higher price. This accords well with the economist’s notion of price discrimination, the notion being that it would be discriminatory not to charge those who create higher costs a higher price. Roughly speaking, if price to cost ratios are the same for each group then discrimination does not occur. However, Hoy and Lambert (2000) show that if an immutable characteristic (such as geno-type) that is related to risk type is used and it is an imperfect signal, then although the more accurate is the information the fewer people are “misclassified”, it is also the case that for those who are misclassified there is a higher the price-cost differential. If one assumes that the impact of discrimination is not linear in price-cost ratios, and is strictly convex instead, then aggregate discrimination may rise as a result of the use of increasingly informative signals.

The view of discrimination often put forward by bioethicists, who are much more active in this research area, is that it is unfair if a “system” – be it public or private – treats some people more harshly due to differences that are beyond their control (such one’s gender or geno-type). Using a genetic test in such a setting leads to price differentials that implement “unfair” discrimination. However, this does not imply that a ban will unambiguously eliminate discrimination. If the Rothschild-Stiglitz separating equilibrium were to obtain as a result of a ban, then the party “discriminated against” (in this case, the high-risk types) would receive no better treatment than if there were no ban. The low-risk types would be worse off and, although they voluntarily choose the policy with a lower level of coverage, one could argue that they were discriminated against in the quantity dimension. Certainly, for the case of life insurance, the survivor families of low risk types tend to end up very badly off as a result of the ban. (See Hoy and Ruse (2008)).

8.4.5 Acquisition of Public Information

In this section we examine incentives for gathering public information. We continue to assume that all consumers initially possess hidden knowledge, having privately observed the outcome of experiment α . Outcomes of the second experiment β , however, are now observed publicly.

Let us first consider the case in which the β -experiment reveals to insurers, but not to consumers, information about the latter’s underlying loss probability. A special case of this environment is considered by Crocker and Snow (1986), where the consumer has already observed the outcome of the α -experiment (α^H or α^L) which is fully informative of the individual’s underlying probability of loss, and in which the β -experiment consists of observing

consumer traits, such as gender, that are imperfectly correlated with the private information held by insurance applicants. The β -experiment provides no information to consumers, who already know their types, but is informative to the informationally constrained insurers. As discussed earlier in section 8.3, this type of categorization, in which the outcome of the β -experiment is publicly observable, enhances efficiency when consumers know *a priori* the outcomes that they will observe for the β -experiment (i.e., their gender). Specifically, a consumer of either β type is at least as well off with categorization based upon β as without it.

Since the β -experiment is not informative for consumers concerning their loss propensities, and does not in any other way influence their preferences, the set of feasible contracts does not depend on whether consumers have prior knowledge of β . Moreover, because each consumer, regardless of β type, is at least as well off with categorization, each consumer must expect to be at least as well off when the outcome of the β -experiment is not privately known *ex ante*. Thus, it is efficient for insurers to categorize applicants on the basis of a publicly observed experiment that is informative for insurers but not for insurance applicants.

The analysis is somewhat different when the β -experiment reveals to consumers information about their underlying loss propensities. In this instance, public information could have a negative social value. As an example, Figure 14 illustrates the extreme situation in which the underlying probability $p^1 = 0$ or $p^2 = 1$ is perfectly revealed by the outcome of the experiment β . Pooling contracts based on β that provide H^* to those revealed to have incurred the loss and A^* to everyone else would allow consumers to attain the same expected utility levels they would realize in the absence of experiment β , when they self-select either H^* or A . Whenever the pair (H^*, A^*) at least breaks even collectively, experiment β has positive social value. It follows that β is socially valuable if and only if the first-best pooling contract lies below the point $F^* \equiv \lambda H^*$

$+ (1 - \lambda)A^*$ in Figure 14. In that event, those consumers revealed to have incurred the loss can be fully compensated by redistributing some of the gains realized by those who have not incurred the loss, permitting attainment of an allocation Pareto superior to (H^*, A) .

When the first-best pooling contract lies above F^* , no redistribution of the gains can fully compensate those revealed to have incurred the loss. In these instances, public information has a negative social value. No insurable risk remains after the public information is revealed, hence its social value is determined by the stronger of two opposing effects, the efficiency gains realized by eliminating adverse selection and the costs of classification risk.¹⁷

As in the case of hidden information, private incentives for gathering public information may not accord with its social value when consumers initially possess hidden knowledge. In the example depicted in Figure 14, the market outcome (H^2, L^1) that occurs when public information is available prior to contracting provides an expected utility equal to the expected utility of the endowment, which is always below the expected utility realized by α^L -types at A and α^H -types at H^* . It follows that, in the present context, the costs of risk classification always discourage the gathering of public information whether or not that information would enhance efficiency.

In contrast with the symmetric information environment, in which public information used to classify consumers has negative social value, when consumers initially possess hidden knowledge, public information can have a positive social value. In the symmetric information environment, the use of public information imposes classification risk on consumers with no countervailing gains in contractual efficiency. However, in markets with asymmetric information, risk classification reduces adverse selection inefficiencies, and these gains may outweigh the costs of classification risk.

8.4.6 Equity/Efficiency Tradeoffs

Concerns about the distributional equity effects of risk classification are not limited to the results of genetic tests. Using gender, age, or race, as well as genetic test results to price insurance coverage may be deemed “unfair” discrimination or otherwise inconsistent with societal norms. As we have seen, these traits would be used in competitive insurance pricing when they are correlated with unobservable risk characteristics. Banning their use to avoid their undesirable distributional equity consequences creates adverse selection inefficiencies. Hoy (2006) and Polborn et al. (2006) refer to the government-created externalities associated with proscriptions on the use of informative risk classification as “regulatory” adverse selection.

To investigate the equity/efficiency tradeoffs involved in public policies concerning the use of risk classification in insurance pricing, consider the case where insurers and insurance applicants are symmetrically informed about each applicant’s risk class, but insurers can be prohibited from using this information in pricing insurance coverage. Implementing such a ban entails foregoing first-best insurance contracting to achieve competing distributional equity goals. One possible approach to analyzing this tradeoff is to quantify separately the distributional and efficiency effects of such a ban. This is the approach taken by Finkelstein et al. discussed in section 8.3.2 above. While this approach has the advantage of making the tradeoff explicit, it has the twin disadvantages of (i) not providing an explicit answer to the question of whether the distributional benefits outweigh the efficiency costs, and (ii) not providing any guidance about the correct way to evaluate the tradeoff between the two quantities.

Hoy (2006) and Polborn et al.(2006) observe that there is often a natural way to strike the balance between distributional equity and allocative efficiency by adopting the “veil-of-

ignorance” [Harsanyi (1953, 1955)], or “contractarian” [Buchanan and Tullock (1962)] methodology to assess the social value of individual utilities. In this approach, although risk class is public knowledge, each consumer’s welfare is evaluated as though the individual were behind a hypothetical veil of ignorance with respect to identity, including risk class. Thus each consumer’s welfare is evaluated as though belonging to the high-risk class (p^H) with probability λ .¹⁸ Further, as consumers are ignorant of their true identities, they adopt the utilitarian social welfare function.

It follows from the observations of section 8.2.2 that a regulation banning the use of risk class in pricing insurance eliminates exposure to classification risk (which is a relevant concern behind the veil of ignorance), but fails to efficiently insure the financial risk because of the induced regulatory adverse selection. Adopting the veil of ignorance approach, Hoy (2006) and Polborn et al. (2006) show that, for some market equilibria, the social benefit of avoiding classification risk can outweigh the social cost of the regulatory adverse selection. Although each analysis investigates a unique environment, the essence of their arguments can be illustrated by relaxing the exclusivity assumption underlying the price-quantity competition that sustains equilibria in the Rothschild-Stiglitz model.

Hoy (2006) observes that, when exclusivity can be practiced, the social cost of regulatory adverse selection always outweighs the social benefit of avoiding classification risk if insurance markets attain the separating Rothschild-Stiglitz equilibrium. Since the market uses risk class to price insurance competitively, utilitarian social welfare in the absence of a ban on its use is given by

$$\lambda V(p^H, H^*) + (1 - \lambda)V(p^L, L^*) \quad (19)$$

as both risk types fully insure at actuarially fair prices, whereas social welfare under a ban on the use of risk class in pricing is the expected value of the Rothschild-Stiglitz contracts,

$$\lambda V(p^H, H^*) + (1 - \lambda)V(p^L, A). \quad (20)$$

Social welfare is clearly lower when the ban is in place, since L -types have lower expected utility in the Rothschild-Stiglitz equilibrium, as they bear the cost of the deductible.

However, when insurers cannot practice exclusivity, applicants are free to purchase from the market any amount of coverage at a price p that is the same for all applicants, and that also results in zero profit given the applicants' coverage choices. Figure 15 illustrates the linear-pricing equilibrium that arises with non-exclusivity. Optimizing the choice of coverage along the same equilibrium opportunity locus, H -types over-insure opting for H , while L -types under-insure, choosing L . Recognizing that these equilibrium choices, along with the equilibrium price of coverage, depend on λ , utilitarian social welfare in the linear-pricing equilibrium can be written as

$$\lambda V(p^H, H(\lambda)) + (1 - \lambda)V(p^L, L(\lambda)), \quad (21)$$

where the contracts resulting in the contingent wealth allocations $H(\lambda)$ and $L(\lambda)$ satisfy the zero-profit condition

$$\lambda[p(\lambda) - p^H]I^H(\lambda) + (1 - \lambda)[p(\lambda) - p^L]I^L(\lambda) = 0, \quad (22)$$

given the coverage $I^t(\lambda)$ optimal for risk class p^t at the market price $p(\lambda)$.

To show that social welfare can be higher when risk classification is banned and insurers cannot practice exclusivity, subtract (19) from (21) and consider the effect of increasing λ starting from $\lambda = 0$, while maintaining the zero-profit condition (22). One obtains

$$\begin{aligned} \frac{\partial}{\partial \lambda} \Big|_{\lambda=0} & \left\{ \lambda [V(p^H, H(\lambda)) - V(p^H, H^*)] + (1 - \lambda) [V(p^L, L(\lambda)) - V(p^L, L^*)] \right\} \\ & = V(p^H, H(0)) - V(p^H, H^*) - U'(W - p^L D)(p^H - p^L) I^H(0). \end{aligned} \quad (23)$$

The third term on the right-hand side of the equality is the marginal effect of an increase in λ on $V(p^L, L(\lambda))$. With $\lambda = 0$, we have $L(0) = L^*$, providing L -types with full-and-fair insurance.

Thus, as an envelope result, the marginal change in their coverage, $\partial I^L / \partial \lambda$, has no effect on social welfare, leaving only the general equilibrium effect of an increase in λ on the price they pay for insurance, as dictated by the zero-profit condition (22).¹⁹

To establish that equation (23) has a positive value in some environments, consider the case of constant absolute risk aversion, where $U(W) = -\exp[-W]$ and the optimal indemnity for an H -type, $I = I^H(0)$, satisfies the first-order condition

$$-p^L(1 - p^H) \exp[-W + p^L I] + (1 - p^L)p^H \exp[-W - (1 - p^L)I + D] = 0.$$

Using this equation, $V(p^H, H(0))$ can be written as

$$\begin{aligned} & -(1 - p^H) \exp[-W + p^L I^H] - p^H \exp[-W - (1 - p^L)I^H + D] \\ & = -(p^H / p^L) \exp[-W - (1 - p^L)I^H + D]. \end{aligned}$$

Hence, equation (23) is positive if

$$\begin{aligned} & 1 - (p^H / p^L) \exp[-(1 - p^L)I^H + (1 - p^H)D] \\ & > (p^H - p^L)I^H \exp[-(p^H - p^L)D], \end{aligned} \quad (24)$$

which is obtained from (23) after dividing by $V(p^H, H^*) = -\exp[-W + p^H D]$. As p^H approaches p^L , the right-hand side of inequality (24) approaches zero, while the left-hand side approaches one. It follows that inequality (24) holds and equation (23) is positive when p^H is

close to p^L . Thus, a ban on the use of information revealing risk class can increase utilitarian social welfare when insurers cannot practice exclusivity and the proportion of H -types is sufficiently low.

Hoy (2006) obtains a stronger result by demonstrating that, regardless of probabilities and the degree of risk aversion, utilitarian social welfare is higher when insurance markets attain a Wilson anticipatory (pooling) equilibrium in the regime where the use of risk class in pricing insurance is banned if the proportion of H -types is sufficiently low. Polborn et al. (2006) derive a similar result in a rich, two-period model of contracting in competitive life insurance markets, where insurers cannot practice exclusivity to deal with regulatory adverse selection.

In each instance, a ban on the use of risk classification in insurance pricing results in allocative inefficiency. Nonetheless, from an ex ante perspective, each consumer trades off gains in prospective H -type utility against losses in prospective L -type utility at the same rate and, when the proportion of H -types is sufficiently small, they agree that the ban increases the expected value of being an H -type by more than it reduces the expected value of being an L -type. Intuitively, the smaller the proportion of H -types, the smaller is their effect on market price, which both mitigates the harm to L -types and enhances the gain to H -types.²⁰

Thus, the veil-of-ignorance approach can place sufficient relative weight on the distributional equity concerns associated with adverse risk classification to endorse bans on the use of public information in some lines of insurance despite their adverse effects on allocative efficiency. Nonetheless, eliminating a ban on the use of such information passes Samuelson's test for potential Pareto improvement.

The resolution of these conflicting prescriptions is illustrated in Figure 16 where, again, H -types choose H and L -types choose L in the linear-pricing equilibrium. The dashed lines

through H and L depict iso-profit lines for H -type and L -type contracts, respectively. The lines must intersect along the pooled fair-odds line, since the zero-profit condition is satisfied. The allocations labeled H' and L' thus also jointly yield zero profit, and are preferred to H and L , respectively, as they provide full coverage at fair marginal prices for the same real cost. Moreover, both allocations can be implemented once risk-based insurance pricing is permitted. It follows that any linear-pricing equilibrium is Pareto dominated by an allocation that categorizes applicants by risk class.

The problem is that, when the redistribution needed to compensate for the adverse equity effects of risk classification is not actually implemented, some non-Paretian ethical judgment must validate eliminating a ban on risk classification, since eliminating the ban results in the pair (H^*, L^*) as the alternative to (H, L) , rather than (H', L') , to the benefit of L -types at the expense of H -types. The hypothetical veil-of-ignorance approach to deriving a social ranking of alternative public policy regimes offers a cogent alternative to the ethical judgments of the Samuelson hypothetical compensation test that we have employed in the preceding sections to address the need for a non-Paretian evaluation of the distributional effects of public policy reforms.

8.5 COMPETITIVE MARKET EQUILIBRIUM AND EXTENSIONS OF THE BASIC MODEL

Although we have emphasized efficiency possibilities in a stylized model of risk classification by insurers, our discussion has practical implications insofar as no critical aspect of insurance contracting is omitted from the model that would have a qualitative effect on efficiency possibilities, and unregulated markets for insurance exploit potential efficiency gains. In this

section, we address the issue of market equilibrium and the implications of several innovations of the model to account for additional features relevant to insurance contracting.

8.5.1 Competitive Market Equilibrium

As shown by Hoy's (1982) original analysis of risk categorization based on immutable characteristics, predictions concerning the performance of an unregulated, competitive insurance market depend on the equilibrium concept employed to account for the presence of asymmetric information. Although the appropriate equilibrium concept remains an unsettled issue, the first empirical evidence was reported by Puelz and Snow (1994) and supported theories that predict the separating Rothschild and Stiglitz allocation (i.e., the pure Nash strategy equilibrium suggested by Rothschild and Stiglitz (1976), the non-Nash reactive equilibrium proposed by Riley (1979) in which insurers anticipate profitable competing entrants, or the take-it-or-leave-it three-stage game analyzed by Cho and Kreps (1987) in which the informed insurance applicants move first), rather than those predicting either a pooling allocation (which can occur in the non-Nash anticipatory equilibrium suggested by Wilson (1977) in which the exit of unprofitable contracts is anticipated, the disassembling equilibrium advanced by Grossman (1979), or the three-stage game analyzed by Hellwig (1987) in which the uninformed insurers move first) or separation with all risk types paying the same constant price per dollar of coverage (as in the linear-pricing equilibrium suggested by Arrow (1970) and analyzed by Pauly (1974) and Schmalensee (1984)).

The evidence reported by Puelz and Snow, however, was also inconsistent with the presence of cross-subsidization between types, first analyzed by Miyazaki (1979) in labor market context, and cross-subsidization is necessary for second-best efficiency in the stylized model unless high-risk types are sufficiently prevalent, as shown by Crocker and Snow (1985a).

Moreover, if competition always leads to the Rothschild and Stiglitz allocation, then the model predicts that the market fails to exploit efficiency gains available through risk categorization based on immutable traits, since all categories have the same risk types represented, so that customers in every category would choose from the same menu consisting of the Rothschild and Stiglitz contracts.

Bond and Crocker (1991) have shown that categorization based on the observed consumption of a product that is correlated with underlying risk alleviates and, in some instances, can eliminate the problem of adverse selection. If endogenous risk classification is imperfect, then further categorization based on immutable traits may be exploited by an unregulated market even in the absence of cross-subsidization when different categories have different risk types represented as a result of the insurer's simultaneous risk classification based on behavior by the insured that influences the risk of loss.

Our discussion of incentives for information gathering reveals that, when categorization is informative for insurance applicants, incentive compatibility constraints are irreversibly altered, and the social value of this type of information could therefore be positive or negative depending on parameters of the environment. As our analysis shows, private incentives for information gathering may not be consistent with efficiency. In unregulated markets, public information or additional hidden knowledge may be acquired when it has negative social value, but go unexploited when it has positive social value.

8.5.2 Extensions of the Model

We have abstracted from a number of considerations that may be of practical relevance to insurance contracting. Here we shall take note of three which appear to be particularly relevant to risk classification.

8.5.2.1 *Multiple Periods.*

Categorization of risks through experience rating is a common practice in insurance contracting, which we have ignored in this review by analyzing an atemporal model. The analysis of Cooper and Hayes (1987) reveals the critical factors that influence contracting with asymmetric information in temporal contexts. For an environment with adverse selection, (costless) experience rating has positive social value if and only if experience is serially correlated with hidden knowledge, as when risk of loss is hidden knowledge and unchanging over time.

The overriding factor determining whether unregulated, competitive markets exploit the efficiency gains of experience rating is the ability of insurers and insurance customers to commit credibly to long-term contracts. If they can, and the market attains the pure strategy Nash equilibrium, then high-risk types receive full and fair insurance, while the coverage and premium for low-risks types is adjusted in the second period based on experience in the first. However, if insurance customers cannot credibly commit to a two-period contract, then experience rating is less valuable as a sorting device, and when renegotiation is introduced, the separating equilibrium degenerates to replications of the single-period equilibrium, as shown by Dionne and Doherty (1991). Hosios and Peters (1989) showed that accident underreporting is possible with commitment and renegotiation, further limiting the market's ability to exploit efficiency gains available through experience rating.

8.5.2.2 *Moral Hazard.*

We have abstracted from moral hazard as a source of informational asymmetry, focusing exclusively on adverse selection. In many insurance markets, however, both informational asymmetries influence contracting possibilities and, as shown by Cramér (1990), the pure strategy Nash equilibrium can be strongly affected by the presence of an unobservable action

taken by the insured that influences the risk of loss. In some instances, moral hazard eliminates the adverse selection problem, and thereby eliminates any social value to risk categorization. In other instances, moral hazard constitutes a new source of nonexistence of a pure strategy Nash equilibrium, and the social value of risk categorization may be enhanced if risk types can be grouped in categories for which the Nash equilibrium exists.

8.5.2.3 *Risk Preferences.*

In the stylized model, all insurance applicants have the same preferences for risk bearing, giving rise to a single crossing of indifference curves for applicants of different risk type. In practice, the willingness to bear risk differs among consumers and is also not directly observable by insurers. Smart (2000) shows that incentive compatibility constraints and the market equilibrium can be fundamentally altered when risk preferences as well as risk type are hidden knowledge, since indifference curves of different risk types may cross twice because of differences in the willingness to bear risk.

In some instances, the qualitative properties of the incentive constraints and the pure strategy Nash equilibrium are not affected, but when differences in risk preferences are sufficiently great, the pure strategy equilibrium, if it exists, may entail pooling of different risk classes, which is inefficient relative to separating contracts. Additionally, for some risk preferences firms charge premiums that are actuarially unfair, resulting in partial coverage with strictly positive profit. For these environments, the model is closed by a fixed cost of entry that dissipates profits through excessive entry. In each of these instances, categorization based on observable traits, either immutable or endogenous, that are correlated with willingness to bear risk has the potential to provide insurers with information that reduces the variation in risk aversion within categories sufficiently to avoid the additional adverse selection inefficiencies

created when insurance applicants with hidden knowledge of risk class have different risk preferences.

8.6 SUMMARY AND CONCLUSIONS

In insurance markets with symmetric information, opportunities for risk pooling can be fully exploited so that perfectly competitive market outcomes are first-best efficient, and consumers are charged actuarially fair premia for insurance coverage. In such markets, information and attendant risk classification have negative social value, even when the information is public, because of the classification risk that must be borne by consumers.

For insurance markets with asymmetric information, risk classification enhances efficiency possibilities. Whether effected through self-selection by insurance applicants possessing hidden knowledge of riskiness (signaling by choice of deductible) or through *a priori* categorization by insurers based on observable traits or behaviors correlated with riskiness (gender, age, race, smoking, or driving sporty cars), risk classification provides insurers with information that relaxes the incentive compatibility constraints and mitigates adverse selection inefficiency.

The unambiguous social benefit of permitting insurers to categorize applicants based upon observable characteristics (such as gender, age or race) that are imperfectly correlated with underlying loss probabilities depends crucially on the assumption that such classification is informative to insurers, but not to their customers. When applicants are fully informed of their underlying loss probabilities, the use of risk classification by insurers expands, and in no way diminishes, the set of feasible (incentive compatible) insurance contracts. Put differently, the

pre-categorization insurance contracts are always feasible in the post-categorization regime. It is the nesting of the regimes that guarantees the efficiency of categorical discrimination.

In contrast, when consumers obtain information about their underlying loss probabilities from the classification procedure (such as in the case of a genetic test), the act of categorization immutably and irreversibly alters the feasible set of insurance contracts. The insurance possibilities that were feasible prior to the classification procedure are precluded by the consumers' changed information sets, which alters the incentive constraints faced by the social planner when designing optimal insurance contracts. Since the pre- and post-categorization regimes are not nested when consumers are informed by the classification procedure, such classification has ambiguous social value.

The adverse equity consequences of risk classification are of special concern to policy analysts when information reveals that some consumers are, in fact, uninsurable. As emphasized by Hoy (1989), these concerns are compounded when action could be taken to diminish the severity of loss, but consumers are discouraged from gathering information and taking such action. We have shown that in markets with either symmetric or asymmetric information, private incentives for initially acquiring hidden knowledge accurately reflect its social value. However, in markets with asymmetric information, private incentives for gathering either public information or additional hidden knowledge are not necessarily consistent with the goal of efficiency in insurance contracting.

The adverse equity consequences of risk classification are precisely the effects that underlie the costs of classification risk. Although we have emphasized these costs as the factor responsible for discouraging consumers from gathering information that has positive social

value, we may also observe that these costs appropriately discourage the gathering of information that has negative social value.

8.7 REFERENCES

- Arrow, Kenneth J. (1970). "Political and Economic Evaluation of Social Effects and Externalities", in *The Analysis of Public Output*, Julius Margolis, ed. New York: Columbia University Press (for NBER).
- Bond, Eric W. and Keith J. Crocker. (1991). "Smoking, Skydiving and Knitting: The Endogenous Categorization of Risks in Insurance Markets with Asymmetric Information," *Journal of Political Economy* 99, 177-200.
- Buchanan, James M. and Gordon Tullock. (1962). *The Calculus of Consent*. Ann Arbor: The University of Michigan Press.
- Cho, In-Koo. and David M. Kreps. (1987). "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics* 102, 179-221.
- Cooper, Russell, and Beth Hayes. (1987). "Multi-period Insurance Contracts," *International Journal of Industrial Organization* 5, 211-231.
- Crocker, Keith J. and Arthur Snow. (1985a). "The Efficiency of Competitive Equilibria in Insurance Markets with Asymmetric Information," *Journal of Public Economics* 26, 201-219.
- Crocker, Keith J. and Arthur Snow. (1985b). "A Simple Tax Structure for Competitive Equilibrium and Redistribution in Insurance Markets with Asymmetric Information," *Southern Economic Journal* 51, 1142-1150.

- Crocker, Keith J. and Arthur Snow. (1986). "The Efficiency Effects of Categorical Discrimination in the Insurance Industry," *Journal of Political Economy* 94, 321-344.
- Crocker, Keith J. and Arthur Snow. (1992). "The Social Value of Hidden Information in Adverse Selection Economies," *Journal of Public Economics* 48, 317-347.
- Cromb, I.J. "Competitive Insurance Markets Characterized by Asymmetric Information," Ph.D. thesis, Queens University (1990).
- Dionne, Georges, and Neil A. Doherty. (1991). "Adverse Selection, Commitment and Renegotiation with Application to Insurance Markets," *Journal of Political Economy* 102, 209-235.
- Dionne, Georges and Nathalie Fombaron. (1996). "Non-convexities and The Efficiency of Equilibria in Insurance Markets with Asymmetric Information," *Economics Letters* 52, 31-40.
- Doherty, Neil A. and Lisa Posey. (1998). "On the Value of a Checkup: Adverse Selection, Moral Hazard and the Value of Information," *Journal of Risk and Insurance* 65, 189-212.
- Doherty, Neil A. and Paul D. Thistle. (1996). "Adverse Selection with Endogenous Information in Insurance Markets," *Journal of Public Economics* 63, 83-102.
- Dreze, Jacques H. (1960). "Le paradoxe de l'information", *Economie Appliquee* 13, 71-80; reprinted in *Essays on Economic Decisions Under Uncertainty*, Cambridge University Press: New York (1987).
- Grossman, Herschel I. (1979). "Adverse Selection, Dissembling, and Competitive Equilibrium," *Bell Journal of Economics* 10, 336-343.
- Harris, Milton and Robert M. Townsend. (1981). "Resource Allocation under Asymmetric Information," *Econometrica* 49, 33-64.

- Harsanyi, John C. (1953). "Cardinal Utility in Welfare Economics and in the Theory of Risk Taking," *Journal of Political Economy* 61, 434-435.
- Harsanyi, John C. (1955). "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility," *Journal of Political Economy* 63, 309-321.
- Hellwig, Martin (1987). "Some Recent Developments in the Theory of Competition in Markets with Adverse Selection," *European Economic Review* 31, 391-325.
- Hirshleifer, Jack. (1971). "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review* 61, 561-574.
- Hosios, Arthur J., and Michael Peters. (1989). "Repeated Insurance Contracts with Adverse Selection and Limited Commitment," *Quarterly Journal of Economics* 104, 229-253.
- Hoy, Michael. (1982). "Categorizing Risks in the Insurance Industry," *Quarterly Journal of Economics* 97, 321-336.
- Hoy, Michael. (1989). "The Value of Screening Mechanisms under Alternative Insurance Possibilities," *Journal of Public Economics* 39, 177-206.
- Hoy, Michael. (2006). "Risk Classification and Social Welfare," *The Geneva Papers* 31, 245-269.
- Hoy, Michael and Peter Lambert. (2000). "Genetic Screening and Price Discrimination in Insurance Markets," *Geneva Papers on Risk and Insurance Theory* 25, 103-130.
- Hoy, Michael, Fabienne Orsi, Francois Eisinger and Jean Paul Moatti. (2003), "The Impact of Genetic Testing on Healthcare Insurance," *Geneva Papers on Risk and Insurance : Issues and Practice* 28, 203-221
- Hoy, Michael and Mattias K. Polborn. (2000). "The Value of Genetic Information in the Life Insurance Market," *Journal of Public Economics* 78, 235-252.

- Hoy, Michael and Michael Ruse. (2005). "Regulating Genetic Information in Insurance Markets," *Risk Management and Insurance Review* 8, 211-237.
- Hoy, Michael and Michael Ruse. (2008). "No Solution to This Dilemma Exists: Discrimination, Insurance, and the Human Genome Project," University of Guelph Discussion Paper No. 2008-8.
- Hoy, Michael and Julia Witt. (2007). "Welfare Effects of Banning Genetic Information in the Life Insurance Market: The Case of BRCA 1/2 Genes," *Journal of Risk and Insurance* 74, 523-546.
- Marang-van de Mheen, P. J., M. C. Maarle and M. E. A. Stouthard. (2002). "Getting Insurance after Genetic Screening on Familial Hypercholesterolaemia; The Need to Educate Both Insurers and the Public to Increase Adherence to National Guidelines in the Netherlands," *J. Epidemiol Community Health* 56, 145-147.
- McDonald, Angus S. (2009). "Genetic Factors in Life Insurance: Actuarial Basis," *Encyclopedia of Life Science (ELS)*, John Wiley and Sons, Ltd: Chichester. DOI: 10.1002/9780470015902.a.0005207.pub2
- Milgrom, Paul and Nancy Stokey. (1982). "Information, Trade and Common Knowledge," *Journal of Economic Theory* 26, 17-27.
- Miyazaki, Hajime. (1977). "The Rat Race and Internal Labor Markets," *The Bell Journal of Economics* 8, 394-418.
- Myerson, Roger B. (1979). "Incentive Compatibility and the Bargaining Problem," *Econometrica* 47, 61-73.

- Oster, Emily, Ira Shoulson, Kimberly Quaid, and E. Ray Dorsey. (2010). "Genetic Adverse Selection: Evidence from Long-Term Care Insurance and Huntington Disease," *Journal of Public Economics* 94, 1041-1050.
- Pauly, Mark V. (1974). "Overinsurance and Public Provision of Insurance: The Roles of Moral Hazard and Adverse Selection," *Quarterly Journal of Economics* 88, 44-62.
- Polborn, Mattias K., Michael Hoy, and Asha Sadanand. (2006). "Advantageous Effects of Regulatory Adverse Selection in the Life Insurance Market," *Economic Journal* 116, 327-354.
- Puelz, Robert and Arthur Snow. (1994). "Evidence on Adverse Selection: Equilibrium Signaling and Cross-Subsidization in the Insurance Market," *Journal of Political Economy* 102, 236-257.
- Riley, John G. (1979). "Informational Equilibrium," *Econometrica* 47, 331-359.
- Rothschild, Casey. (2011). "The Efficiency of Categorical Discrimination in Insurance Markets," *The Journal of Risk and Insurance* 78, 267-285.
- Rothschild, Michael and Joseph E. Stiglitz. (1976). "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics* 90, 630-649.
- Samuelson, Paul A. (1950). "Evaluation of Real National Income," *Oxford Economic Papers* N.S. 2, 1-29.
- Schmalensee, Richard. (1984). "Imperfect Information and the Equitability of Competitive Prices," *Quarterly Journal of Economics* 99, 441-460.
- Smart, Michael. (2000). "Competitive Insurance Markets with Two Unobservables," *International Economic Review* 41, 153-169.

Spence, Michael. (1978). "Product Differentiation and Performance in Insurance Markets,"

Journal of Public Economics 10, 427-447.

Tabarrok, Alexander. (1994). "Genetic Testing: An Economic and Contractarian Analysis,"

Journal of Health Economics 13, 75-91.

Wilson, Charles A. (1977). "A Model of Insurance Markets with Incomplete Information,"

Journal of Economic Theory 16, 167-207.

Footnotes

¹ See Crocker and Snow (1986) for references to U. S. Supreme Court rulings disallowing gender-based categorization in pensions, and to discussions of the laws and public policies related to categorization practices. Tabarrok (1994) provides further references to the policy and popular debate on categorical discrimination.

² Even though the shape of the locus FA is ambiguous, concavity is guaranteed around F. Indeed, the slope of this locus (see Crocker and Snow (1986) page 448) is the right-hand side of condition (c) evaluated at $\delta = 0$: $\frac{\lambda(1-p^H)U'(W_1^L) + (1-\lambda)(1-p^L)U'(W_2^H)}{\lambda p^H U'(W_2^L) + (1-\lambda)p^L U'(W_2^H)}$. Since we have

$W_1^H = W_2^H = W_1^L = W_2^L$ at F, the slope can be rewritten as follows: $\frac{\lambda(1-p^H) + (1-\lambda)(1-p^L)}{\lambda p^H + (1-\lambda)p^L}$.

This reduces to $\frac{1-\bar{p}}{\bar{p}}$, which is the slope of the aggregate zero-profit line. So the AF locus is tangent to the aggregate zero-profit line (see Dionne and Fombaron (1996)).

³ This nomenclature arises because this is the particular allocation supported as an equilibrium in the analyses of Miyazaki (1977), Wilson (1977) and Spence (1978).

⁴ Figure 4 depicts the portion of the utilities possibilities frontier that is better for L -types than the pooling contract F. As discussed in Crocker and Snow (1985a), there is a symmetric portion of the frontier above the 45° line that is better for H -types.

⁵ Both Harris and Townsend (1981) and Myerson (1979) have demonstrated that no alternative organization of the economy's allocation process can dominate the allocations attainable by a social planner.

⁶ So, for example, in the efficiency problem just considered, the goal of the social planner is to maximize the expected utility of one arbitrarily selected agent (V^L) subject to the constraints of (i) not making the other agent worse off than a specified level of expected utility

\bar{V}^H ($V^H \geq \bar{V}^H$); (ii) the economy's resource constraint (5); and (iii) the informational constraints of the market participants (6). By varying \bar{V}^H , the entire set of (second-best) efficient allocations may be determined.

⁷ Since Hoy was concerned with comparing equilibrium allocations in the pre- and post-categorization regimes, the pertinent efficiency issue – can be the winners from categorization compensate, in principle, the losers – was not considered. As Crocker and Snow (1986) demonstrate, the answer to this question, at least in the case of the Miyazaki equilibrium, is that they can.

⁸ An actual Pareto improvement requires that at least one type of agent be made better off while no agents are made worse off. A potential Pareto improvement requires only that the winners from the regime change be able, in principle, to compensate the losers, so that the latter would be made no worse off from the move. As Crocker and Snow (1985b) have demonstrated, there exists a balanced-budget system of taxes and subsidies that can be applied by a government constrained by the same informational asymmetries as the market participants, and which can transform any potential Pareto improvement into an actual improvement. In the discussion that follows, we will use the term “Pareto improvement” to mean “potential Pareto improvement”, recognizing throughout that any potential improvements can be implemented as actual improvements.

⁹ Since the expected utility of an uninformed agent is $\lambda V^H + (1 - \lambda)V^L$ where V^i represents the agent's utility in the informational state i , the slope of the associated indifference curve is $dV^H/dV^L = -(1 - \lambda)/\lambda$.

¹⁰ The Rothschild and Stiglitz allocation is the Pareto dominant member of the class of *informationally consistent* allocations, which is defined as the set of contracts that satisfy self-selection, and that each make zero profit given the class of customers electing to purchase them. While the analysis of the previous sections indicate that these allocations are not always elements of the efficient set (for some parameter configurations), we will, in the interests of expositional ease, assume that they are in the arguments that follow. This is without loss of generality, for in cases where cross-subsidization between risk types is required for efficiency, the same arguments will apply, except with the zero-profit loci relabeled to effect the desired level of subsidy.

¹¹ The problem arises because the H-types have no insurable risks when $p^H = 1$. Whenever $p^H \neq 1$, the allocations B and L depicted in Figure 6 are non-degenerate (in the sense that they do not correspond with the origin). This holds even when $p^L = 0$, although in this particular case the allocation L would reside on the horizontal axis. In contrast, when $p^H = 1$, B and L necessarily correspond with the origin, so there are no insurance opportunities for the uninformed agent (since B is degenerate). This argument holds for any $p^L \geq 0$.

¹² For example, the expected utility of α^L -types is given by

$P(\beta^2|\alpha^L)V(p^2, \hat{H}) + P(\beta^1|\alpha^L)V(p^1, A)$, where the allocation \hat{H} is depicted in Figure 11 below. Using the self-selection condition $V(p^2, \hat{H}) = V(p^2, A)$, we can rewrite this expression as $P(\beta^2|\alpha^L)V(p^2, A) + P(\beta^1|\alpha^L)V(p^1, A)$, which is equal to $V(p^L, A)$ since $P(\beta^2|\alpha^L)p^2 + P(\beta^1|\alpha^L)p^1 = p^L$. Thus, the pair (\hat{H}, A) provides α^L -types the same expected utility that they enjoy at A .

¹³ These profits could then be rebated to the consumers through lower premiums, so that they would be made strictly better off in the post β -experiment regime.

¹⁴ By construction in Figure 10, the α^L -types are indifferent between A , and observing the β -experiment followed by a selection of H^2 or A^1 .

¹⁵ Marang-van de Meheen, van Maarle and Stouthard (2002).

¹⁶ See Joly, Braker, Le Huynh (2010) for details. See also Hoy and Ruse (2005) for a discussion of the broader issues.

¹⁷ The result of Crocker and Snow (1992) showing that public information always has positive social value applies in a linear signaling environment with risk neutral consumers, so the classification risk has no social cost.

¹⁸ In environments where risk class is not known by consumers, as in section 8.2.2, the veil of ignorance is an actual veil with respect to risk class, leading to the same measure of consumer welfare.

¹⁹ Differentiating the zero-profit condition (22) with respect to λ and evaluating the result with $\lambda = 0$, while recognizing that $p(0) = p^L$ and $I^L(0) = D < I^H(0)$ yields $\partial p(\lambda) / \partial \lambda |_{\lambda=0} = (p^H - p^L)[I^H(0) / D]$. Hence, the premium increases by $(p^H - p^L)I^H(0)$.

²⁰ Hoy and Polborn (2000) obtain a yet stronger result showing that, when some consumers are uninformed demanders in the life insurance market, social welfare can increase when they become informed. From an ex ante perspective, uninformed consumers gain from the opportunity to purchase insurance knowing the risk class to which they belong in a manner similar to the analysis in section 8.4.2. Further, in the linear-pricing equilibrium, newly informed demanders may be less risky than the average of those initially in the market, in which case the equilibrium price declines to the benefit of all demanders.




Figure 1








Figure 2





Utilities Possibilities Frontiers: Costless Categorization

Figure 5



Utilities Possibilities Frontiers: Costly Categorization

Figure 6




Figure 7

Relative Risk Aversion	Required Per-Person Endowment Needed to Achieve Utility Level from Non-Categorizing Equilibrium When Categorization is Allowed								Redistribution to Women (\bar{R}^W), Per Woman, % of Endowment		Efficiency Cost Per Dollar of Redistn
	Women (E^W)		Men (E^M)		Total Population (E)		Efficiency Cost as % of Total Endowment				
	MWS (1)	SS (2)	MWS (3)	SS (4)	MWS (5)	SS (6)	MWS (7)	SS (8)			
$\gamma=1$	1.020	1.071	0.979	0.929	0.9998	1	0.038%	0%	2.08%	7.14	3.66%
$\gamma=3$	1.033	1.071	0.966	0.929	0.9998	1	0.025	0	3.39	7.14	1.45
$\gamma=5$	1.040	1.071	0.959	0.929	0.9998	1	0.018	0	4.06	7.14	0.89

Notes: Estimates are based on the model and algorithm described in the text. Columns labeled MWS refer to the high efficiency cost/low redistribution end of the range of possible consequences which obtain when the market implements the Miyazaki-Wilson-Spence equilibrium when gender-based pricing is banned. Columns labeled SS refer to the zero efficiency cost/high redistribution end of the range which obtain when the market implements a pooled-fair full insurance "social security-like" outcome when gender-based pricing is banned. The MWS contracts are computed using Equation (6) and the risk type distributions estimated in Table 1, pooled across genders. Columns (1)-(6) are computed using Equation (15) and columns (9)-(10) are computed using (16).

Table 1




Figure 8




Figure 9




Figure 10




Figure 11




Figure 12




Figure 13




Figure 14

Family Background	Prob of BC (next 10 years)	Prob of BRCA mutation	Prob of BC given BRCA Positive	Prob of BC given BRCA Negative
Low-risk	0.013	0.001	0.141	0.012
High-risk	0.029	0.065	0.295	0.011

Table 2




Figure 15




Figure 16