

Risk Classification in Insurance Contracting

Georges Dionne

3000, Cote-Ste-Catherine, room 4454
Montreal (Qc) Canada, H3T 2A7
Ph. 514.340.6596
Fax 514.340.5019
georges.dionne@hec.ca

Casey G. Rothschild

Pendleton Hall East, Rm. 414
Wellesley College
106 College St.
Wellesley, MA 02481
Ph. 781.283.3026
Fax 781.283.2177
crotsch@wellesley.edu

24 April 2012

Abstract

Risk classification refers to the use of observable characteristics by insurers to group individuals with similar expected claims, compute the corresponding premiums, and thereby reduce asymmetric information. Risk classification can be used to mitigate adverse selection and improve insurance market efficiency, but it may have undesirable equity or efficiency consequences. We employ a canonical screening model of insurance contracting to study these trade-offs in a range of informational environments, and to understand when efficiency or equity concerns are likely to be particularly important. We also review empirical studies on risk classification and residual asymmetric information.

Keywords:

Adverse selection, Classification risk, Diagnostic test, Empirical test of asymmetric information, Financial equity, Insurance rating, Insurance pricing, Moral hazard, Risk classification, Risk characteristic, Risk pooling, Risk separation, Social equity.

1. Introduction

Risk classification refers to the use of observable characteristics such as gender, race, behavior, or the outcome of genetic tests to price or structure insurance policies. Risk classification helps insurers classify selected risks when underwriting. It allows them to group individual risks with similar expected medical costs, compute the corresponding insurance premiums, and reduce adverse selection (and potential moral hazard). Only risk characteristics correlated with expected claim costs are useful for underwriting. Information on individual risk is seldom used to determine individual participation in employer- or government-sponsored plans, but is often observed in voluntary plans and in long-term insurance markets, where it serves to define accessibility, classify policyholders in homogenous risk classes, and set the premiums of each risk class.

In the U.S., the standards of risk classification in the insurer underwriting policy are identified in the Actuarial Standards of Practice 12 (ASOP 12). The choice of rating variables consists of three steps: 1) selection of objective risk characteristics correlated to medical expenses (a causal relationship is not necessary); 2) use of acceptable data and statistical methods that may, however, be subject to the actuary's judgment; and 3) consideration of new or additional characteristics, taking into account the benefits and costs of doing more detailed analysis (American Academy of Actuaries, 2001).

An efficient risk classification system, using actuarial rules and principles, should generate an actuarial premium that reflects the expected cost associated with a given risk characteristic. Two clients with the same risk level should pay the same premium. This is known as the "financial equity" criteria in the actuarial literature (Baker, 2011). In health insurance, premiums are most commonly determined by age, sex, and smoking behavior. Current medical conditions (high cholesterol, diabetes, etc.) and medical histories of older clients are often added as criteria because they can affect the medical expenses covered by the insurance plan. Information on lifestyle, diet, and exercise can also be added.

Risk classification may take into account advances in diagnostics and treatments. However, this practice is limited by regulation in many countries where some characteristics or tests cannot be used to evaluate eligibility and establish premiums. Opinions differ about the use of sophisticated diagnostic tests, such as genetic tests (Chiappori, 2006; Hoy and Witt, 2007). Although these tests clearly provide useful information, they reduce high-risk individuals' access to medical treatments. Often, they reveal medical conditions that are exogenous to the individuals, and considering them for insurance

accessibility raises equity concerns. (This is sometimes referred to as “classification risk.”) This explains the presence of legal constraints on the application of certain tests for insurance classification and pricing. The new Patient Protection and Affordable Care Act (PPACA), introduced in the U.S. in 2010, forbids pricing on any characteristics other than age, family status, geographic or rating area, and tobacco use. The Act restricts even the use of these characteristics (Baker, 2011; Harrington, 2010a, 2010b), and prohibits the pricing of insurance based on individual health status.

More generally, selection and pricing activity based on individual characteristics is subject to concerns about social fairness (or equity) and potential discrimination. This is particularly true in the medical and disability insurance markets. The policy decision to restrict the use of risk classification thus often involves a trade-off between financial and social equity. Less obviously, it also often involves a trade-off between social equity and efficiency provision of insurance.

For example, risk-pooling arising from legal restrictions on risk classification may lead to a situation in which lower-risk individuals pay higher premiums than those corresponding to their true risk, while higher-risk individuals pay lower premiums. On the one hand, these financial inequities (may) reduce classification risk and the potential for discrimination, and improve social equity. On the other hand, when the proportion of high-risk individuals is significant, premiums may rise considerably because low-risk individuals will leave the insurer. The net result can be that higher-risk individuals are no better off as a result of the legal restrictions, and lower-risk individuals, who no longer avail themselves of insurance opportunities, are made strictly worse off.

This situation corresponds to a pure adverse selection “death spiral” problem. One way to reduce this sort of premium spiral is to offer menus of contracts within the risk classes. For example, the new Act in the U.S. permits the sale of deductibles to people under 30, which will probably create a separation of risks within this group. Compulsory insurance is another way to mitigate adverse selection and the need for risk classification.

In a recent paper, Einav and Finkelstein (2011) develop a graphical method that can be adapted to study the costs and benefits of risk classification. Their approach uses a very specific “deadweight-loss” notion of efficiency and does not directly address distributional concerns. It also abstracts from explicitly modeling the nature of insurance market contracting, and, in particular, the screening models that have been the hallmark of the insurance market literature since Rothschild and Stiglitz (1976). In this paper, we survey and elaborate on this “screening” literature. We show how canonical insurance contracting

frameworks can be used to illustrate explicitly the trade-offs between efficiency and distributional concerns, and, moreover, can be used to identify settings in which the equity (“distributional”) or efficiency consequences are theoretically most likely to be particularly important. These theoretical models and consequences are the focus of sections 2-7. Section 8 relates these models to empirical studies on asymmetric information and risk classification.

2. The Canonical Modeling Framework

2.1 Basic Framework

Our formal analysis of risk classification employs a standard adverse selection model of insurance markets. We also consider some moral-hazard extensions in Section 7. Individuals in this model are endowed with a lump sum of money W and face a risk of a monetary loss of size D . Individuals differ in their probability of experiencing this loss. Any two individuals with the same loss probability are said to have the same “risk type”. There is some set I of possible risk types, and p^i is the loss probability of type $i \in I$. For example, $I = \{H, L\}$, with $p^H > p^L$ (a high risk type H and a low risk type L) in the canonical model of Rothschild and Stiglitz (1976).

Insurance contracts are offered by risk neutral insurers. Contracts consist of a premium R paid by the insured and an indemnity M which is paid if the insured suffers a loss. An individual who purchases such a contract thus has net wealth $C_1 = W - R$ available for consumption if she does not experience a loss and $C_2 = W - D + M - R$ if she does experience a loss. It is frequently more convenient to describe contracts directly in terms of the induced consumption allocations (C_1, C_2) rather than premium-indemnity pairs; we follow this convention here.

To illustrate the connection between the premium-indemnity and the induced-consumption descriptions of insurance contracts, Figure 1 below depicts a representative insurance contract (C_1^*, C_2^*) in consumption space. This is the consumption induced by the premium R and the net indemnity $M - R$. This contract (C_1^*, C_2^*) lies below the 45° line, indicating that the indemnity M is less than the loss D . This can be interpreted as a contract with a positive deductible. A contract with a positive deductible is said to provide less than full insurance. A “full insurance contract” has no deductible, and is associated with a consumption point (C_1^*, C_2^*) that lies on the 45° line. (Consumption allocations above the 45° line result from contracts providing “overinsurance.”)

<Figure 1 about here>

An insurer earns profits:

$$\begin{aligned}\pi(C_1, C_2, p^i) &= \left((1 - p^i)W + p^i(W - D) \right) - \left((1 - p^i)C_1 + p^iC_2 \right) \\ &= R - p^iM\end{aligned}\quad (1)$$

from selling a contract (R, M) which induces consumption (C_1, C_2) to a type i individual. Curves of constant profit (iso-profit curves) are thus lines with constant slope $-(1 - p^i)/p^i$ when drawn in consumption space.

Individual are expected utility maximizers, with preferences given by:

$$V(C_1, C_2, p^i) = (1 - p^i)u(C_1) + p^iu(C_2), \quad (2)$$

where $u'(C) > 0$ and $u''(C) < 0$.

Indifference curves have slope $-(1 - p^i)u'(C_1)/p^iu'(C_2)$. Since $u'' < 0$, this means that they are bowed towards the origin (in consumption space), and they are tangent to iso-profit lines at (and only at) points on the full insurance line.

2.2 Informational Environment

Information about risk type can be symmetric or asymmetric. If it is symmetric—i.e., if insurers have the same information about any given individual's risk type as the individual does—then, without loss of generality, we can assume each individual knows her own risk type and that insurers perfectly observe the risk type of each potential customer. (This is because preferences and profits are linear in probabilities: a model with symmetrically known but uncertain risks p^i is analytically identical to a model with symmetrically known and certain risks $\tilde{p}^i = E(p^i)$.)

We will focus attention on asymmetric information environments (or imperfect risk classification based on observable characteristics) in which the insurer is less informed than any given individuals about accident risk inside risk classes. (See Villeneuve, 2005, for an analysis of markets in which the insurer is more informed than the insured.) Moreover, we assume, without loss of generality, that individuals are fully informed about their risk type. If insurers have no information about individuals' risk, then all individuals are indistinguishable—from an expected risk perspective—to insurers: they know nothing

more than that any given individual is drawn from the population distribution of risk types. We use $\Lambda(i)$ to denote this population distribution. Alternatively, insurers may be partially informed: partially informed insurers observe an informative “signal” $\sigma \in \Sigma$ about each individual’s risk type. The signal is informative insofar as the conditional probability distribution $\Lambda(i|\sigma)$ over risk types differs from $\Lambda(i)$. Signals are also referred to in the literature as “groups”, “categories” or “classes”.

When there are two risk types H and L , we can take $\Lambda(H) = \lambda$ and $\Lambda(L) = 1 - \lambda$ so that λ measures the population fraction of high risk types. Similarly, when there are two groups $\sigma \in \{A, B\}$ —say males and females—we can take $\Lambda(H|A) = \lambda_A$ and $\Lambda(H|B) = \lambda_B > \lambda_A$. This yields the model studied by Hoy (1982) and Crocker and Snow (1986), which plays a central role in our analysis. In this case, group A individuals (e.g., men) are less likely to be high-risk individuals, but if $\lambda_A > 0$, there are nevertheless some high-risk individuals in group A . Similarly, group B individuals (e.g., women) are more likely to be high-risk types, but if $\lambda_B < 1$, there are still low-risk types in group B .

2.3 Market Outcomes

Of particular interest is the profile of consumption allocations (C_1^i, C_2^i) which each risk type i obtains in a given market environment. These are frequently (but not exclusively) referred to as “equilibrium” outcomes in the literature, but the literature is somewhat inconsistent in the formal underpinnings of various equilibrium concepts (see Hellwig, 1987, for a discussion) and we therefore employ the more general term “market outcomes” here.

The market outcome which obtains in a given market environment naturally depends on the informational and the institutional features of that environment. We focus on two types of institution, both of which are competitive and both of which involve exclusive contracting. We refer to these two types as the “fixed contracts” and “screening” cases.

In the fixed contracts case, firms compete on price (premium) to provide an exogenously fixed type of contract. This is a useful stylized model for examining settings in which, by law or by custom, insurance contract are standardized, in the sense that coverage (i.e., the losses covered and the indemnities provided for those losses) is uniform across insurance providers. Our focal example of the fixed contracts case involves full-insurance contracts—i.e., contracts fully indemnify the loss D . We assume that the market outcomes in this case are given by the Nash-equilibrium of a game in which (a large number of firms) firms first set contract prices—potentially differing for individuals with different risk

types or in different groups—and individuals then choose the lowest priced contract available to their risk type and/or group.

In the screening case, firms compete on two dimensions: on price and on the type or level of coverage provided. This additional flexibility allows firms to offer menus of contracts designed to differentially appeal to different risk types. A single firm might, for example, offer both a high-deductible low-premium contract and a low-deductible high-premium contract in the hopes of inducing individuals to reveal their risk type via their contract choices. Such a screening strategy is potentially useful when firms cannot observe risk type or face regulatory restrictions on offering type-specific contracts.

The canonical market outcome in the screening case is the Rothschild-Stiglitz (1976) equilibrium, depicted in Figure 2 for a market with two risk types and unobservable private information. This market outcome allocates the contract \vec{C}^{H*} to H -types and \vec{C}^{L*} to L -types. The former lies at the intersection of the 45° line and the H -type's zero-profit line (the line through the no-insurance consumption point $(W, W - D)$ with slope $-(1 - p^H)/p^H$). That is, it provides full insurance at an actuarially fair rate. The L -type's contract \vec{C}^{L*} lies at the intersection of the L -type's zero-profit line and the H -type's indifference curve through \vec{C}^{H*} . It thus provides actuarially fair but incomplete insurance for the lower risk L -types. Specifically, L -types get as much insurance as possible at their actual fair rate without inducing H -types to switch to the contract designed for L -types. Screening of different risk types is achieved in this market outcome because the lower risk L -types strictly prefer \vec{C}^{L*} to \vec{C}^{H*} (as shown by their indifference curve, labeled V^L , through \vec{C}^{L*}) and are willing to sacrifice full insurance to credibly signal that they are lower risks and thereby receive lower premiums.

Rothschild and Stiglitz (1976) show that the market outcome $(\vec{C}^{H*}, \vec{C}^{L*})$ depicted in Figure 2 is the only possible Nash equilibrium of a game in which multiple firms simultaneously offer individual contracts and then individuals choose their preferred contract from the set of contracts offered. They also show that there is no equilibrium of the game when the fraction λ of high risk types is sufficiently low. This lack of equilibrium occurs because when λ is low, a portion of the pooled actuarially fair line (i.e., the zero-profit line for an individual with risk type $\bar{p} = \lambda p^H + (1 - \lambda)p^L$) lies above the L -type's indifference curve. When this is the case, a firm could deviate from the Rothschild-Stiglitz candidate equilibrium by offering a single contract providing consumption between the pooled fair line and the L -type indifference curve. This would be a profitable deviation since it would attract both risk types; the Rothschild-Stiglitz equilibrium candidate is thus not, in fact, a Nash equilibrium in this case.

<Figure 2 about here>

The literature has primarily employed three distinct alternative equilibrium concepts to resolve this non-existence problem. Each of the three concepts predicts a unique equilibrium outcome for any set of parameters, but they do not always coincide, and there is no consensus about which one (if any) is right (see Hellwig, 1987). The Riley (1979) reactive equilibrium predicts that the Rothschild-Stiglitz candidate equilibrium will be the market equilibrium, even when the proportion of high risk types is low. The Wilson (1977) (E2) foresight equilibrium coincides with the Rothschild-Stiglitz equilibrium whenever the latter exists, and otherwise is a pooling equilibrium in which both risk types receive the contract on the pooled zero-profit line which is most preferred by the L -type.

The Miyazaki (1977)-Wilson (1976)-Spence (1978) (henceforth, MWS) involves a similar foresight concept and recognizes that at the Rothschild-Stiglitz candidate equilibrium, there may be scope for Pareto improving cross subsidies from the L -types to the H -types. The contract pair $(\vec{C}^{H*}, \vec{C}^{L*})$ in Figure 3 illustrates an example of such a cross-subsidy. The L -type contract \vec{C}^{L*} lies strictly below the L -type's zero-profit line, while the H -type contract \vec{C}^{H*} lies strictly above the H -type's zero profit line. The pair thus involves cross subsidies. Since the point \vec{C}^{H*} still provides full insurance, it is obviously strictly better for the H -types than the Rothschild-Stiglitz candidate contract, labeled $\vec{C}^{H\#}$. As drawn, the L -type contract \vec{C}^{L*} also lies on a higher indifference curve (labeled V^L) than L -types would be on in the Rothschild-Stiglitz candidate equilibrium (the dashed indifference curve below V^L). This is possible because the cross-subsidy to H -types makes H -types less interested in switching to the L -type's contract. This allows the market to provide more insurance to L -types. Pareto improving cross-subsidies exist whenever the benefits to L -types of this additional insurance are more than enough to outweigh the costs of providing the cross-subsidy.

<Figure 3 about here>

In the MWS equilibrium, the market takes advantage of these Pareto-improving cross subsidies whenever they exist. Specifically, the MWS equilibrium is the member of the class of constrained efficient separating allocations which maximizes the expected utility of L -types subject to H -types being at least as well off as they would be with their full insurance actuarially fair contract. For sufficiently high values of λ , it again coincides with the Rothschild-Stiglitz equilibrium. For lower values of λ , it involves cross-subsidies from low-risks to high-risks. As in the Rothschild-Stiglitz candidate equilibrium, H -types

still receive full insurance, but now at potentially better-than-actuarially fair rates. Similarly, L -types remain underinsured—albeit less so—but pay greater than actuarially fair premiums. (See Dubey and Geanakoplos, 2002, Martin, 2007, Netzer and Scheuer, 2008, and Picard, 2009 for modern discussions of the justifications of the Rothschild-Stiglitz and MWS concepts.)

An alternative approach to modeling market outcomes in the insurance literature is to focus on efficient outcomes (see, e.g., Crocker and Snow, 2000) rather than equilibrium outcomes *per se*—although the two can coincide, since the MWS equilibrium is constrained efficient (for a fixed proportion of high-risk types; see Crocker and Snow, 1985a). A “first best efficient” market outcome maximizes aggregate profits:

$$\sum_{i \in I} \Lambda(i) \pi(\vec{C}^i, p^i) \quad (3)$$

subject to

$$V(\vec{C}^i, p^i) \geq \bar{V}^i \quad (4)$$

for all risk-types $i \in I$, where \bar{V}^i is a minimum utility constraint for the i type. It is straightforward to show that an allocation is first best efficient for some set of minimum utilities if and only if $C_1^i = C_2^i$ for each i . A social planner who observes each individual’s risk types and can dictate their required insurance contracts could, in principle, implement any first-best efficient market outcome with non-negative aggregate profits. Since we typically assume that the social planner is interested in the well-being of individuals and not firms *per se*, we focus on the first best efficient market outcomes with exactly zero profits. (Indeed, the term first best efficient frequently incorporates this zero profit assumption, as in Crocker and Snow, 2000.)

When individuals are privately informed about risk type and the social planner is completely uninformed about an individual’s risk type a “second-best” efficiency problem for the social planner is appropriate. In this second-best (or “constrained”) efficient problem, the social planner faces the additional incentive compatibility constraints:

$$V(\vec{C}^i, p^i) \geq V(\vec{C}^j, p^i) \quad \forall i, j. \quad (5)$$

These constraints reflect the fact that the social planner cannot directly identify an individual's type, and thus cannot force an individual to select a contract which he finds strictly worse than some other contract which is available.

When the social planner is partially informed about individual risk types via a partially informative signal σ , it can condition contracts on the signal, so the incentive compatibility constraints for the constrained efficient problem are instead:

$$V(\vec{C}^{i,\sigma}, p^i) \geq V(\vec{C}^{j,\sigma}, p^i) \quad \forall i, j, \sigma \text{ such that } \Lambda(i|\sigma) > 0, \Lambda(j|\sigma) > 0 \quad (6)$$

Our approach to studying the effects of risk classification on market outcomes will be to compare market outcomes in the presence of risk classification to market outcomes in its absence.

It is worth noting that we do not consider some other models of market outcomes which arise in the literature. For example, we do not explicitly consider models with a monopoly insurance provider (e.g., Stiglitz, 1977) or oligopolistic markets (e.g., Buzzacchi and Valletti, 2005). Nor do we consider linear pricing equilibrium, which are often used in markets where contracting is non-exclusive and individuals can buy small amount of coverage from multiple providers simultaneously (e.g., Pauly, 1974, Hoy and Polborn, 2000, and Villeneuve, 2003).

3. Risk Classification can have Purely Distributional Consequences

Policy discussions about risk classification frequently emphasize the perceived distributional benefits of restricting firms from employing risk classification and downplay or ignore the potential efficiency costs of such restrictions. We first consider a (rather restrictive) setting in which risk-classification does not have any efficiency consequences and this emphasis is appropriate.

This setting is characterized by the following assumptions:

- (i) there are **two risk types**, H and L , with $p^H > p^L$, and a fraction λ of H -types;
- (ii) there is **symmetric information**, so insurance providers can readily identify the risk type of any individual;
- (iii) insurance contracts are **fixed full insurance contracts**, so firms compete only on price;

- (iv) there is a **mandatory purchase** requirement: each individual must buy one (and only one) insurance contract.

Figure 4 depicts the market outcomes. When insurance providers classify based on observable risk type, each type $i \in \{H, L\}$ pays its type-fair premium $R^i = p^i D$ in exchange for full indemnification of the loss D . The associated consumption allocations are labeled \vec{C}^{i*} . When insurance providers do not employ risk classification based on risk type, both types of individuals pay the pooled-fair premium $\bar{R} = (\lambda p^H + (1 - \lambda)p^L)D \equiv \bar{p}D$, with results in consumption \vec{C}^{P*} . This is the unique equilibrium outcome of the game in which firms first post prices (which are type specific if risk classification is employed) and individuals then choose from the contracts available to them.

<Figure 4 about here>

Risk classification clearly has distributional consequences in this setting: it improves the well-being of the L -types at the expense of the well-being of H -types. Because both market outcomes are first-best efficient, risk classification has no efficiency consequences.

Whenever risk classification is permitted, it will be used by the market. This is because, at the no classification equilibrium (\vec{C}^{P*} in Figure 4) a deviating firm could make positive profits offering a slightly less expensive contract (just up and to the right from \vec{C}^{P*}) which is available only to L -types. This will attract all L -types, leaving only the H -types buying \vec{C}^{P*} and rendering that contract unprofitable. As such, the no-classification market outcome is only likely to arise in the presence of explicit legal restrictions on the use of risk classification. The consequences of risk classification are therefore best interpreted as the consequences of legalizing versus banning risk classification here.

In this simple setting, legal restrictions on pricing based on risk-related categorical affiliations like gender can be imposed without adverse efficiency consequences. This does not mean *ipso facto* that such restrictions—which transfer resources from low-risk individuals to high-risk individuals—are desirable, however. From an *interim* perspective, when risk types are taken as given and known, the desirability of such bans depends on distributional preferences. An objective observer’s preference is likely to depend on why individuals have different risk types. For example, if the riskiness of a given individual is the result of previous health investment decisions such as diet or exercise, then a restriction on the use of observable correlates of this past behavior in pricing policies will have the perverse effect of punishing

good past decisions and rewarding bad ones. An objective observer might reasonably regard this as inequitable.

Even when risk type is completely outside of individual control, however—for example when type is determined by genetic predispositions towards ill health—opinions about the desirability of the distributional consequences of banning risk classification may differ. Low-risk individuals are likely to object to the redistribution that results from such a ban, and might reasonably wonder why they should be forced to effectively subsidize other individuals' insurance premiums. A utilitarian social planner, on the other hand, will always prefer the market outcome with a ban on risk-categorization. As Hoy (2006) points out, utilitarian social welfare preferences are a natural choice: they follow directly from Harsanyi's (1953, 1955) veil of ignorance approach since they coincide with the *ex-ante* preferences of an individual who does not yet know which risk type she will end up being.

The assumptions (i)-(iv) underlying this analysis are obviously quite restrictive. It is straightforward to show that assumptions (i) and (ii) can be relaxed, though: the same basic result would hold if there were more risk types, or if insurers could observe an imperfect signal of individuals' privately known risk type. On the other hand, although the full insurance part can be relaxed, the fixed contract part of assumption (iii) is essential. As we discuss in the following section, if contracts are not fixed, there will typically be negative efficiency consequences of imposing legal restrictions on risk classification in screening settings.

The mandatory purchase assumption (iv) may or may not be essential, depending on specifics of individual preferences and on the distribution of risk types. Figure 4 depicts a situation in which it is not essential. As drawn, the point \vec{C}^{P*} lies strictly above the "autarkic" indifference curves (labeled V^H and V^L) that pass through the no-insurance consumption point $(W, W - D)$. This means that, if risk classification were banned, both individuals would still *voluntarily* choose to purchase the contract \vec{C}^{P*} instead of foregoing insurance. On the other hand, if $V(\vec{C}^{P*}, p^L) < V(W, W - D, p^L)$ so that the L -type indifference curve instead hit the 45° line above \vec{C}^{P*} , then assumption (iv) is essential. We discuss this voluntary purchase case in the following section.

4. Risk Classification Bans can be Purely Efficiency Reducing

The example above may be what some policy makers have in mind when contemplating a ban on risk classification: they imagine that the only result of the imposition of a ban will be a shift to pooled prices and they view this as distributionally desirable. In many settings, this is an unreasonable expectation.

Consider the consequences of relaxing the mandated purchase assumption (iv) in the fixed contracts setting of the preceding section, for example. Without this mandate, a ban on risk-classification can reduce efficiency and have no beneficial distributional consequences at all: If the fraction λ of high-risk types is sufficiently high, then $V(\vec{C}^{P^*}, p^L) < V(W, W - D, p^L)$. In contrast with Figure 4, this means that the L -type indifference curve through the no-insurance point $(W, W - D)$ will hit the 45° line above \vec{C}^{P^*} . If risk classification is banned in such an environment, then any premium at which L -types will buy insurance will be strictly unprofitable when it is also sold to H -types. Such a ban will therefore lead L -types to exit the insurance market; they will be strictly worse off as a result. Moreover, once L -types are driven out of the market, H -types will end up with the contract yielding consumption \vec{C}^{H^*} , exactly as they would if risk classification were legal. Banning risk classification in such an environment is therefore purely efficiency reducing: it hurts L -types without benefitting H -types.

In fact, the negative efficiency consequences of a ban on risk-classification in a fixed contract voluntary-purchase environment can be quite dire. Suppose, for example, that the fixed contract provides full insurance, and that there is a continuum of risk types with risk probabilities p^i uniformly distributed on $[0,1]$. For any given premium $R < D$ (i.e., any policy actually offering any insurance), only individuals with

$$p^i \geq q(R) = \frac{u(W) - u(W - R)}{u(W) - u(W - D)} \quad (7)$$

will voluntarily choose to purchase insurance. The break-even premium to sell to this set of individuals is $D(1 + q(R))/2$. When $\frac{D(1+q(R))}{2} > R$ the premium R is strictly lower than premium required to break even on the set of individuals who voluntarily purchase insurance contracts at premium R , and offering contracts at a premium R will lead to losses. It is easy to find examples of utility functions $u(\cdot)$ with the property that $\frac{D(1+q(R))}{2} > R$ for every $R < D$. (E.g., $u(x) = x - kx^2$ with $k < \frac{1}{2w}$.) In this case there is no premium $R < D$ at which firms are willing to sell contracts. The unique market outcome therefore involves no insurance at all.

This pernicious outcome can be understood intuitively as follows: if firms naively priced insurance at the pooled fair premium ($R = 0.5$), the lowest risk individuals would choose not to purchase and the insurance pool would end up being too risky for that premium. Insurers would be forced to raise premiums, inducing some of the least risky remaining individuals to drop their insurance, forcing firms to further raise premiums, and so forth. Depending on the preferences and the distribution of types, this can lead to a complete unraveling of the market. This unraveling was first described formally by Akerlof (1970), and is known as an “adverse selection death spiral.” (See, for example, Cutler and Reber (1998) and Strohmeier and Wambach (2000).)

When contracts are not fixed, firms can typically use screening mechanisms to prevent a complete collapse of the market for insurance. Bans on risk classification can still be purely efficiency reducing even when firms employ screening, however. Suppose, for example, that there are two risk types, that information is symmetric, and that market outcomes are described by the Riley equilibrium (i.e., the Rothschild-Stiglitz candidate equilibrium). Then when risk classification is permitted, individuals will purchase their type-specific full insurance actuarially fair contracts \vec{C}^{i*} as in Figure 4. When it is banned, they will instead purchase the contracts \vec{C}^{i*} depicted in Figure 2. Even though both types still continue to actively participate in the insurance market, the ban is purely efficiency reducing: H -types are unaffected by the ban, but L -types now purchase a contract providing strictly less insurance than before and are strictly worse off.

Buchmueller and DiNardo (2002) look empirically at the consequences of community rating (a ban on risk classification) in the small group and individual health insurance market in New York State. They provide evidence that suggests exactly this sort of an effect: they find no evidence of a significant reduction in the fraction of individuals with insurance, but they do find evidence of a significant shift in towards less generous insurance coverage.

5. Risk Classification Can Involve Trade-offs between Efficiency and Distributional Equity

5.1 Efficiency and Distributional Effects with MWS Market Outcomes

The consequences of risk classification in screening environments are less clear cut under other types of market outcomes. We use the MWS market outcomes concept to illustrate that bans on risk classification can have negative efficiency consequences and beneficial distributional consequences.

Figure 3 can be used to analyze the two-type, symmetric information case. With risk classification, individuals would receive full insurance at their actuarially fair full insurance premium, resulting in the consumptions $\vec{C}^{i\#}$. The contracts \vec{C}^{i*} denote the MWS outcomes without risk classification. As drawn, this market outcome involves cross-subsidies from the L -types to the H -types. Because of these cross subsidies, H -types are strictly better off than they would be with their actuarially fair full insurance contract. L -types are better off than they would be without the cross subsidies—i.e., in a Rothschild-Stiglitz equilibrium—but strictly worse off than they would be in the market outcome with the contract $\vec{C}^{L\#}$ that they would receive with risk classification.

For a social planner who would like to redistribute from L -types to H -types, banning risk classification and moving the market outcomes from $\vec{C}^{i\#}$ to \vec{C}^{i*} has desirable distributional consequences. Such a ban has unambiguously negative efficiency consequences, however: it moves the economy from a first-best efficient allocation to one which is not first best efficient. Indeed, the allocation \vec{C}^{i*} is Pareto dominated by the pair $(\vec{C}^{H*}, \vec{C}^L)$ depicted in Figure 3. The pair $(\vec{C}^{H*}, \vec{C}^L)$ yields aggregate profits of zero, and, since information is symmetric and risk type is observable, it is informationally feasible. A social planner could, in principle, directly implement that strictly preferable outcome. Alternatively, by imposing a properly calibrated tax/subsidy on contracts sold to L/H types, it could induce a decentralized market with legalized risk-classification to offer them..

Crocker and Snow (1986) show that the same conclusions generalize when risk classification is only partially informative. They consider the setting with two risk types, and two partially informative signals (“groups” or “categories”) A and B , where signal B (respectively, A) is associated with a greater (respectively, lower) probability of being a high risk. Formally $\Lambda(H|A) = \lambda_A$ and $\Lambda(H|B) = \lambda_B > \lambda > \lambda_A$, with $\lambda_B < 1$ and $\lambda_A > 0$. It is easiest to see how the results generalize by examining the MWS equilibria with and without risk classification. Without risk classification, the market outcome $(\vec{C}^{H*}, \vec{C}^{L*})$ is simply the MWS outcome for a two risk-type economy with a fraction λ of high-risks. With risk classification, the market outcome provides individuals in group σ ($\sigma \in \{A, B\}$) with the consumptions $(\vec{C}^{H*,\sigma}, \vec{C}^{L*,\sigma})$ associated with the MWS outcome for a two risk-type economy with the group-specific fraction λ_σ of high-risks. There are three possibilities (exhaustive, because if an MWS outcome involves positive cross-

subsidies with some fraction of high risks, the MWS outcome in a market with a smaller fraction of high-risks will as well):

- (1) The A -group market outcomes $(\vec{C}^{H^*,A}, \vec{C}^{L^*,A})$ do not involve cross subsidies from low risks to high risks.
- (2) The A -group market outcomes $(\vec{C}^{H^*,A}, \vec{C}^{L^*,A})$ involve cross subsidies, while the no risk-classification market outcomes $(\vec{C}^{H^*}, \vec{C}^{L^*})$ do not involve cross subsidies.
- (3) The no risk-classification market outcomes $(\vec{C}^{H^*}, \vec{C}^{L^*})$ involve cross subsidies.

In case (1), market outcomes coincide with the Rothschild-Stiglitz equilibrium consumptions both with and without risk classification. In this case, banning risk classification has neither efficiency nor distributional consequences—indeed, it has no effects at all.

In case (2), the market outcome with risk classification Pareto dominates the market outcome without risk classification: Category B individuals receive the same (Rothschild-Stiglitz) contracts with and without risk classification, while both risk types in Category A are strictly better off. In this case, risk classification has only (beneficial) efficiency consequences.

Case (3) is illustrated in Figure 5. Triangles denote the category- and type-specific contracts $\vec{C}^{i^*,\sigma}$ when risk classification is legal, and circles denote the type-specific contracts $\vec{C}^{i^*,\sigma}$ when classification is banned. As it shows, risk classification makes category- A individuals better off (moves them to the higher dashed indifference curves) and category- B individuals worse off (moves them to the lower dashed indifference curves).

Risk classification also has efficiency consequences in case (3) because the market outcome without risk classification is not second-best efficient in this case. The MWS equilibrium is the constrained efficient outcome which maximizes the well-being of L -types. When this equilibrium involves cross subsidies, the magnitude of these cross subsidies is chosen to balance the benefit of easing the incentive constraint with the cost of reduced resources. Lowering the fraction of high-risk types lowers the resource cost of any given amount of incentive constraint easing, and thus tilts the balance towards additional cross subsidies. This means that at $(\vec{C}^{H^*}, \vec{C}^{L^*})$ in Figure 5, additional cross subsidies would be Pareto improving within category A . That is, there is therefore an alternative allocation that could be offered in an incentive-compatible way to the category A individuals that would increase the well-being of both category- A types at the same net actuarial cost for the category. (It is not illustrated in Figure 5, but in it

the H -type contract would lie on the full insurance line above and to the right of \vec{C}^{H*} , and the L -type contract would lie up and to the left of \vec{C}^{L*} .) It is therefore feasible to make the category A individuals strictly better off without changing the allocations to category- B individuals.

<Figure 5 about here>

5.2 Welfare Analysis with Distributional and Efficiency Effects

There are three general approaches to welfare analysis when there are both efficiency and distributional consequences. One approach is to adopt an explicit social welfare function which assigns welfare weights to the different types, as in Hoy (2006). This approach can be viewed as imposing a tacit trade-off between the distributional benefits and the efficiency costs. A second approach is to explicitly and separately quantify the efficiency and distributional consequences; this approach is useful when different policy analysts differ in their perception about the optimal trade-off between distribution and efficiency concerns (*viz* Finkelstein et al., 2009, who apply this approach to the analysis of gender classification in the U.K. compulsory annuity market).

A third approach is to argue that policymakers should focus exclusively on the efficiency consequences. Consider, for example, the market outcome $(\vec{C}^{H*}, \vec{C}^{L*})$ which obtains under a ban on risk classification, as depicted in Figure 5. The fact that this is inefficient *means* that there is some alternative allocation which a social planner faced with the same informational constraints as the market could Pareto improve upon it. Drawing on earlier work (Crocker and Snow, 1985b), Crocker and Snow (1986) show, moreover, that this sort of Pareto improvement can be “decentralized” by imposing a set of contract (and category) specific taxes while legalizing risk classification. Hence, although distributional concerns imply that it is possible for a social planner to prefer the no risk-classification outcome $(\vec{C}^{H*}, \vec{C}^{L*})$ to the market outcomes $(\vec{C}^{H*,A}, \vec{C}^{L*,A})$ and $(\vec{C}^{H*,B}, \vec{C}^{L*,B})$ which obtain with legalized risk classification, such a social planner will always strictly prefer to legalize risk classification while imposing an appropriate set of taxes.

Rothschild (2011) discusses an alternative way to “decentralize” this Pareto improvement: the provision of partial social insurance. Figure 6 illustrates his argument. The MWS outcome with banned categorization provides consumption \vec{C}^{i*} to i -types; as illustrated, these involve cross subsidies from L -

risks to H -risks, so risk classification has both distributional and efficiency consequences—i.e., the market outcome falls under case (3), as described above.

Now suppose the government implements a social insurance policy providing an indemnity of M^S and a (mandatory) premium R^S illustrated in Figure 6. This policy is the pooled-fair policy that transfers the same resources from L -types to H -types as the MWS outcome would. In the presence of such a social insurance policy, individuals would have an effective wealth of $\tilde{W} = W - R^S$ and would face an effective loss of $\tilde{D} = D - M^S$. The MWS outcome with banned risk-classification and this effective endowment-loss pair leads to the same consumptions \vec{C}^{i*} as before. In other words, in the presence of a ban on risk classification, the partial social insurance policy has no effect on the market outcome (when described in terms of realized consumptions). But with the partial social insurance policy in place, these market outcomes are now implemented without market driven cross subsidies from L -risks to H -risks. That is, the MWS equilibrium now falls in case 2 and introducing risk classification is purely efficiency improving.

<Figure 6 about here>

This argument boils down to a simple observation: removing bans on risk classification is potentially problematic only insofar as it “undoes” cross-subsidies provided by the market between the different categories. Socializing the cross-subsidy (e.g., via partial social insurance provision) moots these concerns. The argument is therefore quite general. For example, it is not dependent on the assumption that market outcomes are described by the MWS equilibrium, and it applies even when risk-classification is costly for firms to implement, for example because verifying an individual’s risk class involves performing a costly test.

This means that even when the redistributive benefits of banning risk classification make a ban potentially desirable, imposing a ban is never an optimal way to achieve redistributive goals. This argues for evaluating risk classification exclusively on efficiency grounds—at least whenever more efficient ways of achieving those redistributive goals are, in fact, available.

6. Endogenous Informational Environments

In the preceding analysis, we modeled individuals as having perfect information about their own risk type. This is not essential: what is essential for that analysis is that individuals have a fixed set of information about their risk type, and that it is at least as good as the information insurers have.

In many applications, however, the assumption of static information is unrealistic. Individuals may be able to learn more information about their own risk types through a screening test or medical examination, for example. With the advent and development of genetic screening, analysis of insurance purchases in these sorts of “dynamic” informational environments is increasingly important.

Dynamic informational environments are more challenging to analyze because of the interaction between testing availability, decisions, and insurance opportunities. In particular, an individual’s *decision* to learn more about her risk type will likely to depend on how that information will affect one’s insurance options, and those insurance options will in turn depend on individual’s incentives to learn information. For example, an individual might reasonably worry that taking a genetic test exposes them to *classification risk*: they might worry that a bad test outcome will make them uninsurable or insurable only with a dramatically increased premium. This could be problematic even if firms cannot get access to test outcome information without permission from the customer: since individuals with good test outcomes will have an incentive to voluntarily reveal that outcome to insurers, insurers may be able to *infer* bad outcomes from the omission of good test results.

Doherty and Thistle (1996) and Dionne et al. (2000) provide a comprehensive analysis of these issues (see also Crocker and Snow, 1992). We follow their basic analysis here, but modify it in two ways: first, we consider the implications of risk classification in this context more explicitly. Second, we employ the MWS market outcome concept instead of the Rothschild-Stiglitz concept, and thereby resolve some of the non-existence problems that arise in their analyses.

The basic framework is simple: there are two risk types H and L , as above. There are two possible informational states for individuals: informed individuals know their risk type and uninformed individuals (U -types) do not. Let κ denote the fraction of individuals who are informed; since individuals can choose whether or not to become informed, $\kappa \geq \kappa_0$ is an endogenous quantity here, where $\kappa_0 > 0$ is the (exogenous) initial fraction of informed individuals. The fraction of informed individuals who are H -types, denoted again by λ , is independent of κ . Finally, there is a fixed utility cost $\tau \geq 0$ for individuals to take the test. U -types are uncertain about their true risk type, but if they remain uninformed they behave exactly as would an individual with risk type $p^U = \lambda p^H + (1 - \lambda)p^L$.

There are four cases to consider, which differ in the availability of test status information to insurers (or to a notional social planner). Describe the MWS market outcomes in these cases facilitates the analysis of the welfare consequences of banning risk classification in this environment.

6.1 Symmetric information

If information is completely symmetric, the market outcomes are simple (and are the same under MWS and Rothschild-Stiglitz concepts): each type i ($i = H, L, U$) receives a type-specific actuarially fair full insurance contract $\vec{C}^{i\#}$ (i.e., the actuarially fair contract for an individual with risk p^i). Since becoming informed does not change the expected value of U -types' consumption. Since it introduces welfare reducing uncertainty—in the price of the contract received— U -types have no incentive to become informed, and $\kappa = \kappa_0$.

6.2 Private test results, public information status

Figure 7 depicts the market outcome if insurance firms can observe whether an individual has been tested or not—and can condition their contracts on this observation—but do not observe the outcome of that test. In this case, U -types again get their full-insurance actuarially fair contract $\vec{C}^{U\#}$. H - and L -types get the contracts \vec{C}^{H*} and \vec{C}^{L*} associated with the MWS market outcome for a two risk-type economy, which involve cross subsidies (as drawn).

As in the symmetric information case, uninformed individuals again do *not* have an incentive to become informed in this market outcome.

If information status is known and test results are private but verifiable, then L -types will then have an incentive to verifiably reveal their test results and receive their full-insurance actuarially fair contract $\vec{C}^{L\#}$. Moreover, firms will then infer that informed types who did not reveal their risk types are H -types; they will therefore receive *their* full-insurance actuarially fair contract $\vec{C}^{H\#}$. This case is thus effectively the same as the symmetric information case.

<Figure 7 about here>

6.3 Private information status, verifiable test outcomes

If type and informational status are private, but test status is verifiable, then L -types will again reveal their type and receive full fair insurance. We will solve for the market outcome for the U - and H -types in

two steps. First, we consider the market outcome taking the fraction κ of informed types as given. Second, we consider the incentives to become informed to identify the possible equilibrium κ values.

For any given $0 < \kappa < 1$, the market will implement a standard MWS outcome, with U -types playing the role usually played by L -types. Denote by $\vec{C}^{i\$}(\kappa)$ the contract received by i -types ($i = U, L$). Since the ratio of L - to U -types is $\lambda\kappa/(1 - \kappa)$, higher κ reduces the scope for utility-improving cross subsidies from the U - to the L -types. Consequently, $V(\vec{C}^{U\$}(\kappa), p^U)$ is non-increasing in κ . Moreover, there is some $\bar{\kappa}$ such that there are zero cross subsidies when $\kappa \geq \bar{\kappa}$ (when the MWS outcome coincides with the Rothschild-Stiglitz equilibrium) and positive cross subsidies when $\kappa < \bar{\kappa}$. In the latter case, $V(\vec{C}^{U\$}(\kappa), p^U)$ is strictly decreasing in κ ; in the former, it is independent of κ .

The value of information to U -types is:

$$I(\kappa) = \lambda[V(\vec{C}^{H\$}(\kappa), p^H) - V(\vec{C}^{U\$}(\kappa), p^H)] + (1 - \lambda)[V(\vec{C}^{L\#}, p^L) - V(\vec{C}^{U\$}(\kappa), p^L)]. \quad (8)$$

Since the first term is zero (by the binding incentive compatibility constraint in the MWS outcome):

$$I(\kappa) = (1 - \lambda)[V(\vec{C}^{L\#}, p^L) - V(\vec{C}^{U\$}(\kappa), p^L)]. \quad (9)$$

Figure 8 plots $I(\kappa)$, which jumps up discretely at $\kappa = 1$ since the market will offer no contract to (non-existent) U -types there, and choosing to remain uninformed would force U -types into contract $\vec{C}^{H\#}$.

<Figure 8 about here>

The equilibrium values of κ can be found for various test costs τ (we assume that $\kappa_0 < \bar{\kappa}$). When $\tau < I(\kappa_0) < I(1)$, all individuals want to get tested no matter what other individuals do. The market outcome therefore involves all individuals getting tested and receiving their type-specific actuarially fair full insurance contracts. When $I(1) < \tau$, nobody wants to get tested and the unique equilibrium has $\kappa = \kappa_0$. For intermediate values of τ there are multiple equilibrium values of κ . This multiplicity arises from of a coordination problem: as Figure 8 illustrates, the incentive to become informed increases as more and more individuals become informed. When $I(\bar{\kappa}) < \tau \leq I(1)$, there are two equilibria: either all individuals coordinate on remaining uninformed (so that $\kappa = \kappa_0$) or they coordinate on becoming informed and $\kappa = 1$. When $I(\kappa_0) \leq \tau < I(\bar{\kappa})$, there is, additionally, a third (unstable) equilibrium κ : the unique intermediate value satisfying $I(\kappa) = \tau$. Whenever there are multiple equilibria, they can be Pareto ranked, with lower κ equilibria being (generically) strictly better than higher κ equilibria.

6.4 Purely Private information

When information is purely private, the market outcome can be found by first considering the market outcome for any *given* $\kappa \geq \kappa_0$ and then looking for equilibrium values of κ .

When $0 < \kappa < 1$, there are effectively three unobservably distinct types. Figure 9 provides a qualitative depiction of the MWS market outcome $(\vec{c}^{L'}(\kappa), \vec{c}^{U'}(\kappa), \vec{c}^{H'}(\kappa))$. In this outcome H -types receive full insurance, and there are two binding incentive compatibility constraints:

$$\begin{aligned} V(\vec{c}^{H'}(\kappa), p^H) &= V(\vec{c}^{U'}(\kappa), p^H) \\ V(\vec{c}^{U'}(\kappa), p^U) &= V(\vec{c}^{L'}(\kappa), p^U). \end{aligned} \tag{10}$$

Unlike the Rothschild-Stiglitz equilibrium candidate outcomes, however, the MWS outcome *may* involve cross-subsidies. Figure 9 depicts the case where L -types provide cross subsidies to the U -types (i.e., L -types lie below their actuarially fair line and U -types lie above theirs).

<Figure 9 about here>

The value of information to an U -type is strictly positive in any market outcome with $0 < \kappa < 1$. To wit:

$$\begin{aligned} (1 - \lambda)[V(\vec{c}^{L'}, p^L) - V(\vec{c}^{U'}, p^L)] + \lambda[V(\vec{c}^{H'}, p^H) - V(\vec{c}^{U'}, p^H)] \\ = (1 - \lambda)[V(\vec{c}^{L'}, p^L) - V(\vec{c}^{U'}, p^L)] > 0. \end{aligned} \tag{11}$$

That is, uninformed individuals are indifferent to becoming informed if they turn out to be H -types, but strictly prefer becoming informed if they turn out to be L -types. Denote the value of information at the market outcome associated with an informed fraction κ by $I(\kappa)$. In contrast with the private, verifiable information case discussed above, it is not straightforward to sign $I'(\kappa)$. But it is easy to show that $I(1) = 0$ and that $I(\kappa)$ is continuous. It follows that whenever $\tau > I(\kappa_0)$, there is an equilibrium at κ_0 in which no (additional) individuals will choose to take the test and that the unique equilibrium when $\tau = 0$ and testing is costless has $\kappa = 1$.

For intermediate values of $\tau \in (0, I(\kappa_0))$, neither of these outcomes is an equilibrium: if no individuals got tested, then all would want to, and if all individuals got tested, no one would want to. In contrast with Doherty and Thistle (1996), however, there *is* an equilibrium with some intermediate *fraction* $\kappa \in (\kappa_0, 1)$ of individuals taking the test. This value satisfies $I(\kappa) = \tau$, which holds for some κ by the

continuity of $I(\kappa)$. Hence, for any $\tau > 0$, the market outcome is a three-risk type MWS market outcome like the one depicted in Figure 9.

6.5 The Inefficiency of Risk Classification Bans with Endogenous Information

Risk classification is only relevant when firms potentially have access to some information which they can *use* for classification—e.g., in the informational environments of sections 6.1, 6.2, and 6.3. Banning risk classification in each of these settings effectively makes information purely private, as in section 6.4. We now show that such a ban leads to inefficiencies in each of the three settings. To do so, we explicitly construct an incentive and information compatible contract menu that could, in principle, be offered by a social planner with no more information than insurance firms and that would Pareto improve on the market outcome that obtains without risk classification.

Suppose first that when risk classification is banned, the market outcome involves $\kappa < 1$, so there are some informed individuals. The market outcome $(\vec{C}^L, \vec{C}^U, \vec{C}^H)$ is depicted in Figure 10. (Figure 10 remains agnostic about the cross-subsidies.) Figure 10 also depicts an alternative menu of contracts $(\vec{C}_\varepsilon^L, \vec{C}_\varepsilon^U, \vec{C}_\varepsilon^H)$. This menu is constructed in three steps:

1. $C_{\varepsilon,1}^{H'} = C_{\varepsilon,2}^{H'} = C_1^{H'} + \varepsilon = C_1^{H'} + \varepsilon$, so H -types are offered slightly more generous full insurance coverage.
2. $V(\vec{C}_\varepsilon^{U'}, p^H) = V(\vec{C}_\varepsilon^{H'}, p^H)$ and $\pi(\vec{C}_\varepsilon^{U'}, p^U) = \pi(\vec{C}^{U'}, p^U)$, so that H -types remain indifferent to their contract and the U -type's contract, and U -types get additional insurance, and the additional insurance is priced at their actuarially fair rate.
3. $V(\vec{C}_\varepsilon^{L'}, p^L) - V(\vec{C}_\varepsilon^{U'}, p^L) = V(\vec{C}^{L'}, p^L) - V(\vec{C}^{U'}, p^L)$ and $C_{\varepsilon,1}^{L'} = C_{\varepsilon,2}^{L'}$, so that L -types are offered full insurance, and the benefit to a U -type of becoming informed and learning that she is an L -type is unchanged.

<Figure 10 about here>

The allocation $(\vec{C}_\varepsilon^L, \vec{C}_\varepsilon^U, \vec{C}_\varepsilon^H)$ is incentive and information compatible whenever L is observable or verifiable. Moreover, by steps 2 and 3, the value of information to U -types is unchanged; it therefore remains an equilibrium for the same fraction κ of individuals to be informed.

Since $\pi(\vec{C}_0^L, p^L) > \pi(\vec{C}^L, p^L)$, the menu $(\vec{C}_\varepsilon^L, \vec{C}_\varepsilon^{U'}, \vec{C}_\varepsilon^{H'})$ therefore earns non-negative profits for sufficiently small ε . The menu $(\vec{C}_\varepsilon^L, \vec{C}_\varepsilon^{U'}, \vec{C}_\varepsilon^{H'})$ could therefore be implemented by a social planner no more informed than the market. Since it Pareto dominates the banned risk-classification market outcome $(\vec{C}^L, \vec{C}^{U'}, \vec{C}^{H'})$, risk classification is inefficient whenever L is observable or verifiable—i.e., in the informational environments of sections 6.1 and 6.3. When information *status* but not risk type is observable, as in section 6.2, a similar construction can be used to show that $(\vec{C}^L, \vec{C}^{U'}, \vec{C}^{H'})$ is inefficient.

Suppose next that the market outcome when risk classification is banned involves $\kappa = 1$. The market outcome will therefore be the MWS outcome $(\vec{C}^{L*}, \vec{C}^{H*})$ for a two-type economy. This is inefficient in the informational environments of section 6.1 and 6.3, since a social planner can feasibly implement the contract pair $(\vec{C}^{L**}, \vec{C}^{H*})$, where \vec{C}^{L**} is the full insurance contract with $\pi(\vec{C}^{L**}, p^L) = \pi(\vec{C}^{L*}, p^L)$. It is also inefficient in the informational environment of section 6.2. In this case, a social planner can instead offer the menu $(\vec{C}^{L*}, \vec{C}^{U**}, \vec{C}^{H*})$, where \vec{C}^{U**} is the full insurance contract satisfying $\pi(\vec{C}^{U**}, p^U) = \lambda\pi(\vec{C}^{H*}, p^H) + (1 - \lambda)\pi(\vec{C}^{L*}, p^L)$. Uninformed types will strictly prefer \vec{C}^{U**} to becoming informed and facing a lottery (with the same expected value) over the other two contracts, and moreover will not have to bear the cost of getting tested.

We conclude that the outcomes without risk classification are inefficient in each of the three information environments in which risk classification could potentially be used.

6.6 Distributional Consequences of Risk Classification

Risk classification may have distributional as well as efficiency consequences. For example, H -types may receive positive cross subsidies in the pure private information market outcome depicted in Figure 9. When it does, banning risk classification will increase the well-being of H -types relative to the symmetric information market outcomes $\vec{C}^{i\#}$ (*viz* Figure 7).

The welfare consequences of risk classification with endogenous private information are thus qualitatively quite similar to the consequences with fixed private information discussed in section 4. A ban on risk classification may lead to a market outcome which is preferable—on distributional grounds—to the market outcome which obtains when risk classification is permitted, but imposing a ban

is never an optimal way to achieve redistributive goals: the fact that they are inefficient means that there is always be a feasible—though perhaps not practical—menu of Pareto dominant contracts.

6.7 Risk Classification when Testing Is Intrinsically Beneficial

The preceding analysis focused on a situation in which there are no intrinsic benefits of testing; the private benefits of testing, if any, arise from the informational asymmetries and the associated sorting of individuals into insurance contracts. In practice, of course, testing may have direct benefits. It might facilitate better treatment, for example. A simple way of incorporating this sort of benefit is to allow the cost of testing τ to be negative. While this can change the market outcomes discussed in sections 5.1 and 5.2 (since individuals could now choose to get tested), it does not substantively affect the analysis of risk classification in these environments. Hoel et al. (2006) reach a similar conclusion in a model with heterogeneity in the benefits of testing. (Also see Doherty and Posey, 1998, and Hoel and Iversen, 2000 who build on Hoy’s 1989 model of self-protection to consider the consequences of genetic testing in a market with asymmetric information.)

7. Endogenous Risk Classification

We have focused on the pure adverse selection case in the preceding sections, and abstracted from “moral hazard.” *Moral hazard* refers to a situation in which an individual’s riskiness p^i depends on some choice x she makes, and this choice is in turn influenced by her insurance coverage. For example, diet and exercise influence health and hence health risk, and the presence of insurance coverage for statins, beta blockers, and more invasive and expensive interventions like cardiac bypass surgery may lead individuals to be less attentive to their diet and exercise regimens.

The theoretical literature on risk classification in moral hazard contexts is thin. Bond and Crocker (1991) is an exception. They consider a special case of a “moral hazard” framework where individual preferences are given by

$$V^i(C_1, C_2, x) = (1 - p^i(x))u(C_1) + p^i(x)u(C_2) + \theta^i G(x). \quad (12)$$

An individual’s risk thus depends on some decision x , and the decision x has a direct effect on individual well-being via the strictly concave and increasing function $G(x)$. The parameter θ^i captures the taste for

x , which has unit cost c . The accident risk $p^i(x)$ of the two types $i = H, L$ potentially differ for two reasons. First, for any given x , it may be that $p^H(x) > p^L(x)$. Second, different tastes $\theta^H > \theta^L$ may lead different types to make different x choices. For expositional simplicity, we focus here on the second effect by taking $p^H(x) = p^L(x) \equiv p(x)$.

The choice variable x could, for example, represent the extent to which an individual is over-weight, which is (at least in part) a result of individual choices. H -types have a stronger taste for being overweight, either because they enjoy weight-promoting activities like eating rich foods more than L -types, or because their cost of maintaining lower weight is higher (e.g., they dislike weight-reducing activities, or have a metabolism which makes it harder for them to lower their weight). The simplifying assumption that $p^H(x) = p^L(x) \equiv p(x)$ means that individuals are otherwise identical: if they chose the same weight, then they would face the same statistical risk of a loss.

Risk classification involves pricing policies based on the observable decisions x . This is called “endogenous risk classification” because the risk classification depends not on an exogenously determined variable like gender, but rather on the endogenous choice of x . Following Bond and Crocker (1991), it is straightforward to show that endogenous risk classification leads to a first-best efficient market outcome: A first best allocation maximizes $V^i(C_1^i, C_2^i, x^i)$ subject to a resource constraint

$$[p(x)C_2^i + (1 - p(x))C_1^i + cx] - [p(x)(W - D) + (1 - p(x))W] \leq Y^i \quad (13)$$

for each i and for some type-specific resource levels Y^i (e.g., zero for a first best allocation which is break even for type i). Letting λ^i be the Lagrange multiplier on the resource constraint, the first order necessary conditions for C_1^i and C_2^i at a first best allocation are:

$$(1 - p(x))u'(C_1^i) = \lambda(1 - p(x)), \quad (14)$$

and

$$p(x)u'(C_1^i) = \lambda p(x), \quad (15)$$

from which it follows immediately that $C_1^i = C_2^i \equiv C^{i*}$ and $\lambda = u'(C^*)$ in any first best allocation. The first order necessary condition for x^i is

$$p'(x)[u(C_1^i) - u(C_2^i)] + \theta^i G'(x) = \lambda[c + Dp'(x) + p'(x)(C_2^i - C_1^i)], \quad (16)$$

or

$$\frac{\theta^i G'(x^{i*})}{u'(C^{i*})} = c + Dp'(x^{i*}). \quad (17)$$

<Figure 11 about here>

Figure 11 plots the first-best (x^{i*}, C^{i*}) pairs for $i = H, L$, associated with $Y^L = Y^H = 0$. Because $p(x)$ is the same for both types, the zero-profit line labeled $\pi = 0$ is the same for both types. (It is non-linear because of the x -dependence of $p(x)$.) The first best allocations occur at the tangency points of the indifference curves (labeled V^H and V^L), and the tangency point for H -types has $x^{H*} > x^{L*}$ because $\theta^H > \theta^L$. It is clear from Figure 11 that this allocation is incentive compatible. It is therefore the unique Nash equilibrium of the Rothschild-Stiglitz style game in which many profit-maximizing firms simultaneously offer contracts and individuals then choose their optimal contract. This means that endogenous classification—here captured by the fact that firms are offering different contracts to individuals who make different choices of x —leads to a first-best efficient outcome. This conclusion is robust to small differences between $p^H(\cdot)$ and $p^L(\cdot)$, which manifest themselves in Figure 11 as vertical differences between the break-even constraints of the two types.

If classification is banned, so that contracts are not allowed to condition on x , on the other hand, no first best allocation is feasible. To see this, note that utility an individual gets from choosing x and a contract providing an indemnity M at a premium R is:

$$\widehat{U}(x) \equiv p(x)u(W - R - D - cx - M) + (1 - p(x))u(W - R - cx) + \theta^i G(x). \quad (18)$$

At a first best allocation, $C_1 = W - R - D - cx - M = W - R - cx = C_2 = C^*$, and Equation (17) can be used to sign the derivative of $\widehat{U}(x)$:

$$\widehat{U}'(x^*) = \theta^i G'(x) - cu'(C^*) = u'(C^*)[Dp'(x^*)] > 0. \quad (19)$$

This means that at any first-best allocation each individuals have an incentive to choose a higher x . As such, no first-best efficient allocation is implementable when firms cannot employ endogenous risk classification and their contracts on x . Banning endogenous risk classification is therefore inefficient in this simple setting. This inefficiency is intuitive: banning the use of x in pricing policies effectively makes x unobservable and introduces an entirely standard “moral hazard” problem.

8. Risk Classification and Residual Asymmetric Information in Insurance Markets

The objective of this section is to propose different tests for the presence of residual asymmetric information in insurance markets. From the preceding sections, we know that the potential presence of asymmetric information between clients and insurers regarding individual risks motivates risk classification. Full efficiency in risk classification should separate individual risks and generate different actuarial insurance premiums that reflect these risks. This means there should not be any residual asymmetric information between the insurer and the insured inside the risk classes. With actuarial premiums, full insurance should be the optimal contract and there should not be any correlation between insurance coverage and individual risk. But in the real life of health insurance contracting, there are numerous constraints that limit efficiency in risk classification. Incentive contracting becomes important and the empirical question is to what extent these constraints limit the efficiency of risk classification. In the following sections, we present the statistical methodology developed over the recent years to verify if there remains residual asymmetric information in observed risk classes.

Information problems are common in insurance markets. Usually, and as assumed in sections 2-7 above, the insured are better informed about their own characteristics or actions than are their insurer. The two best-known information problems discussed in the economics literature are moral hazard and adverse selection (Arrow, 1963). Asymmetric learning is another information problem that can degenerate into adverse selection over time. (See Cohen, 2005; Dionne et al, 2010.) Symmetric learning, in contrast, can foster create a full information situation between the parties. Doing statistical tests on the link between risk classification and the presence of a given asymmetric information problem is therefore very complicated, because the same correlation between a contract characteristic and an observed risk can be attributed to more than one information problem. It may also be attributed to other characteristics that are not well controlled in the statistical test. The theoretical predictions must be carefully established in a theoretical model (Dionne et al., 2010) to distinguish the effect of each information problem.

Many theoretical contributions were published in the 1970s to account for stylized facts observed in insurance markets. The one-period or static models discussed above were first developed. Partial insurance, such as deductible and co-insurance contracts, can be justified by asymmetric information. However, a deductible can be optimal for moral hazard, adverse selection, or proportional

administrative costs. Risk classification based on observable characteristics is another mechanism associated with the presence of asymmetric information. A separating equilibrium is a form of risk classification in which different risks choose different contracts. As documented above, insureds who accept a higher deductible against a lower premium choose to be in a different risk class than the higher-risk individuals in the portfolio.

Therefore, a crucial empirical question is: Do we need additional risk separation schemes when an efficient risk classification system exists? In other words, does an insurer that applies risk classification techniques based on observable characteristics in its underwriting policy need to search for optimal contract forms within the risk classes? The answer is no in static or one-period contracts, provided there is no residual asymmetric information within the risk classes. Yet the reality is much more complicated because contract duration between the parties can cover many periods, and over time, the true risks may become known to both parties.

The first goal of empirical research on information problems in markets is to determine whether residual asymmetric information remains in these markets. The empirical question in insurance can be summarized as follows: Is there any residual correlation between chosen insurance coverage and risk within risk classes? The second goal is to identify which information problem remains when the first test rejects the null hypothesis that there is no residual information problem. This step is important for the insurer because it must implement the appropriate instruments to improve resource allocation. A deductible efficiently reduces ex-ante moral hazard, but not necessarily ex-post moral hazard because often, the accident has already occurred when the action is taken. A high deductible can even have an adverse effect and encourage accident cost building (Dionne and Gagné, 2001). As is well known in the empirical literature, a positive correlation between insurance coverage and risk is a necessary condition for the presence of asymmetric residual information, but it does not shed light on the nature of the information problem. The third goal is to find ways to improve the contracts and reduce the negative impact of asymmetric information on resource allocation. These resource allocation objectives must take into account other issues such as risk aversion, fairness, and accessibility of services. This last issue is particularly important in health care markets. A decrease in insurance coverage may reduce ex-ante moral hazard because it exposes the insured person to risk, but it also significantly reduces accessibility to health services for sick people who are not responsible for their condition.

In insurance markets, the distinction between moral hazard and adverse selection boils down to a question of causality (Chiappori, 2000). The theoretical literature on moral hazard states that

unobserved actions of the insured result from the forms of contracts. For example, a generous health insurance plan can reduce the incentives for prevention and increase the risk of becoming sick. With adverse selection, the nature of the risk already exists, but the nature of the contracts chosen is a function of the risks. There is therefore a reverse causality between the risk and the contract when we consider each problem separately, although the correlation between insurance coverage and the level of risk is positive in both cases.

Econometricians analyze two types of information when studying insurers' data (Gouriéroux, 1999; Puelz and Snow, 1994; Dionne et al., 2006; Richaudeau, 1999). The first type contains variables that are observable by both parties to the insurance contract. Risk classification variables are one example. Econometricians/insurers combine these variables to create risk classes when estimating accident distributions. They can be used to make estimates conditional on the risk classes or inside the risk classes. The second type is related to what is not observed by the insurer or the econometrician during contract negotiations, but can explain the insured's choice of contracts or actions. If we limit our interpretation to asymmetric information (either moral hazard or adverse selection), we can test the conditional residual presence of asymmetric information in an insurer's portfolio by testing for a correlation between the contract coverage and the realization of the risk variable during a contract period. Two types of tests have been proposed in the literature (Chiappori and Salanié, 2000; Dionne et al., 2001; see Chiappori, 2000; and Chiappori and Salanié, 2003 for detailed analyses). One parametric test estimates the following relationship:

$$y_i = \alpha + \beta X_i + \gamma d_i + \delta E(d_i | X_i) + \varepsilon_i, \quad (20)$$

where y_i is the contract choice by individual i (level of deductible, for example), X_i is a vector of control variables such as the observable characteristics used in risk classification and control variables for risk aversion, β is a vector of parameters to be estimated, d_i is the realization of the random variable observed at the end of the contract period (accident or not, for example), $E(d_i | X_i)$ is the conditional expected value of the random variable obtained from the estimation of the accident distribution, and ε_i is the residual of the regression. A positive sign is usually anticipated for the coefficient of d_i when residual asymmetric information remains (higher coverage or lower deductible is related to more accidents or higher risk). The seminal theories of Rothschild and Stiglitz (1976) and Wilson (1977) strongly predict that such a correlation should be observed in the data in the presence of adverse selection, while Holmstrom (1979) and Shavell (1979) strongly predict that the correlation is due to moral hazard. Note that the dependent variable in the above regression can be the risk variable d_i while

the coverage y_i is an independent variable. This symmetry is discussed in detail in Dionne et al. (2006). The presence of the variable d_i is not necessarily exogenous in equation (20). It is often better to instrument this variable (See Dionne et al., 2009, for more details).

The presence of $E(d_i|X_i)$ is necessary to control for specification errors (missing variables) or for potential non-linearity not modeled in the equation. Without this control, the coefficient of d_i can be significant for reasons other than the presence of residual asymmetric information in the risk classes.

If the coefficient of d_i is not significant, one can reject the presence of residual asymmetric information in the risk classes when all other factors are well controlled. This does not mean that there is no asymmetric information in this market; rather, it means that the insurer's risk classification system eliminates asymmetric information efficiently, and that there is no *residual* asymmetric information within the risk classes. In other words, when risk classification is done properly, it is not necessary to choose the contract form within the risk classes to reduce asymmetric information.

An equivalent model was proposed by Chiappori and Salanié (2000). Here, two equations are estimated simultaneously, one for contract choice and the other for accident distribution. An example is the bivariate probit model:

$$y_i = f(X_i, \beta) + \varepsilon_i \quad (21)$$

$$d_i = g(X_i, \beta) + \eta_i \quad (22)$$

The test consists in verifying whether there is dependence between the residuals of the two equations. An absence of conditional correlation is interpreted as an absence of residual asymmetric information in the data. The authors present an additional non-parametric test that is independent on the functional forms of the above models. It is based on a Chi-square test of independence.

To separate moral hazard from adverse selection, econometricians need a supplementary step. An additional market relationship can be estimated to look for adverse selection (conditional on the fact that the null hypothesis of no asymmetric information was rejected), as Dionne et al. (2009) did for auctions. In insurance markets, dynamic data are often available. Time adds an additional degree of freedom to test for asymmetric information (Dionne and Doherty, 1994; Hendel and Lizzeri, 2003). This information can be used in many insurance markets where past experience is important. Experience rating works at two levels in insurance. Past accidents implicitly reflect unobservable characteristics of the insured (adverse selection) and introduce additional incentives for prevention (moral hazard).

Experience rating can therefore directly mitigate problems of adverse selection and moral hazard, which often hinder risk allocation in the insurance market.

Experience rating not only provides additional information on risk, but may also play an important role in the dynamic relationship between policyholders' insurance claims and contract choice. The theoretical literature on repeated insurance contracting over time clearly indicates that these features may help overcome problems of moral hazard when risks known to the policyholder (endogenous) are unobservable by the insurer (Winter, 2000) or when exogenous characteristics are unobservable (Dionne et al. 2000). Contract choice is influenced by the evolution of the premium, which is closely linked to the insured's risk or past experience. Because increased insurance coverage tends to lower the expected cost of accidents for the insured, incentives for safe behaviour are weakened for all risks. Under experience rating, the subsequent rise in accidents increases the marginal costs of future accidents. Experience rating may therefore offset the disincentive effect created by single-period insurance coverage.

However, the above tests are conducted in a static framework, which fails to recognize the dynamics that experience rating introduces in contractual relationships. Chiappori and Salanié (2000) discuss in detail how the omission of the experience-rating variable, even in tests with one-period data, must plausibly explain the failure to detect asymmetric information. Dionne et al. (2010) tested this conjecture by adding a bonus-malus variable in equations similar to those presented above. They affirmed that the bonus-malus coefficient was indeed negatively related to the level of insurance coverage (through fluctuations in the premium) and positively correlated to claims (potentially through unobserved heterogeneity). The coefficients thus appear to hide the link between claims and contract choice, which is exactly what Chiappori and Salanié (2000) argue. This is apparent in traditional cross-sectional tests, as well as in extrapolations using longitudinal data models that simply pool repeated observations or permit the correlation of unobserved independent factors with each contract observed over time. The additional time factor thus improves the power of the test to detect asymmetric information.

Abbring et al. (2003a, 2003b) apply a multi-period incentive mechanism by focusing on the dynamics of claims, but not on the dynamics of contract choice (because of data limitations). Applying specific assumptions about the wealth effects of accidents to policyholders who differ only in their claim records (thus their experience rating), their model predicts that subjects with the worst claims records should try harder to increase safety, and thereby, *ceteris paribus*, file fewer claims in the future. However, their

data do not support the presence of moral hazard. Dionne et al. (2011) extend their model and do not reject the presence of moral hazard, using a different data set. The potential presence of adverse selection in their data was not a real problem because all drivers must be insured (see also Abbring et al., 2008 for another test of moral hazard).

Dionne, Michaud, and Dachour (2010) show that failure to detect residual asymmetric information, and more specifically, moral hazard and adverse selection in insurance data, is due to the failure of previous econometric approaches to model the dynamic relationship between contract choice and claims adequately and simultaneously when looking at experience rating. Intuitively, because there are at least two potential information problems in the data, an additional relationship to the correlation between risk and insurance coverage is necessary to test for the causality between risk and insurance coverage. Using a unique longitudinal survey of policyholders from France, they propose a methodology to disentangle the historical pathways in claims and premiums. They show how causality tests can be used to differentiate moral hazard from asymmetric learning (and eventually adverse selection). They do not reject moral hazard for a given group of policyholders, and do not reject asymmetric learning for younger drivers.

Cohen and Siegelman (2010) present a survey of adverse selection in insurance markets. They review the empirical studies in different insurance markets, including health insurance and long-term care insurance. They argue that the coverage-risk correlation is particular to each market. Accordingly, the presence of a significant coverage-risk correlation has different meanings in different markets, and even in different risk pools in a given market, depending on the type of the insured service, the participants' characteristics, institutional factors, and regulation. This means that when testing for the presence of residual asymmetric information, one must control for these factors as well.

For example, one reason for not observing a significant correlation between coverage and risk is the "absence of insured private information" on the insured's health status. As suggested by Cohen and Siegleman (2010), the statistical test should be done within the risk classes, even if the employer does not use some observable variables as a risk classification variable. Another reason for the lack of risk-coverage correlation, which may also apply to health insurance, is policyholders' failure to use their private information to negotiate their insurance coverage and premium. Pauly et al. (2003), for example, show that the demand for life insurance is not sensitive to insurance price and risk.

Risk aversion is a private information factor that reduces the empirical link between insurance coverage and risk (Finkelstein and McGarry, 2006; Fang et al, 2008). Cutler, Finkelstein, and McGarry (2008) provide additional evidence for this effect. They verify how insurance coverage, risk, and risk aversion are interrelated in five markets including life, health, and long-term care. They affirm that individuals participating in risky activities (less risk-averse) are less likely to buy insurance. However, these individuals do not represent lower expected claim payments. A possible extension of this research is to use the new proposed cross-sectional tests by Chiappori et al. (2006) based on profit maximization in competitive markets that are robust to heterogeneity in preferences.

As documented by Finkelstein and Porteba (2006), however, insurance consumption depends on institutions. Moreover, risk classification in the health care market is regulated in many countries. Therefore, the empirical predictions based on the implicit assumption of competitive markets may not be appropriate for many markets, including health insurance. For further discussion on particularities other than efficient risk classification that may generate an absence of correlation between insurance coverage and risk, see Cohen and Singleman (2010).

9. Conclusion

In this article we have proposed an economic analysis of risk classification in insurance markets. We have studied in detail the canonical modeling frameworks to illustrate the possible trade-offs between efficient insurance provision and social equity associated with the use of risk classification. We have also analysed different models of risk separation when there are residual informational asymmetries within risk classes. These informational asymmetries can either be fundamental, or can arise from regulatory restrictions on the use of risk-relevant characteristics in pricing insurance policies. Such restrictions are often observed in health insurance because some characteristics or tests are banned from being used by insurers to evaluate different risks. Finally, we have studied the empirical efficiency of risk classification by presenting different empirical tests for the presence of residual asymmetric information in different insurance markets.

References

- Abbring, J., Chiappori, P.A., Heckman, J., and Piquet, J. (2003a). Adverse selection and moral hazard in insurance: Can dynamic data help to distinguish? *Journal of the European Economic Association* **2-3**, 512–521.
- Abbring, J., Chiappori, P.A., and Piquet, J. (2003b). Moral hazard and dynamic insurance data. *Journal of the European Economic Association* **1**, 767–820.
- Abbring, J., Chiappori, P.A., and Zavadil, Z. (2008). Better safe than sorry? Ex ante and ex post moral hazard in dynamic insurance data. Mimeo, VU University of Amsterdam, 69 pages.
- Akerlof, G.A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics* **84**, 488-500.
- American Academy of Actuaries (2001). Risk classification in voluntary individual disability income and long-term care insurance. http://www.actuary.org/pdf/health/issue_genetic_021601.pdf.
- Arrow, K.J. (1963). Uncertainty and the welfare economics of medical care. *The American Economic Review* **53**, 941–973.
- Baker, T. (2011). Health insurance, risk, and responsibility after the patient protection and affordable care act. Research paper 11-03. ILE, University of Pennsylvania Law School.
- Bond, E. and Crocker, K. (1991). Smoking, skydiving, and knitting: The endogenous categorization of risks in insurance markets with asymmetric information. *Journal of Political Economy* **99**, 177-200.
- Brown, J.R. and Finkelstein, A. (2009). The private market for long-term care insurance in the United States: A review of the evidence. *Journal of Risk and Insurance* **76**, 1, 5-29.
- Browne, M.J. and Doeringhaus, H.I. (1993). Information asymmetries and adverse selection in the market for individual medical expense insurance. *Journal of Risk and Insurance* **60**, 300–312.
- Buchmueller, T. and DiNardo, J. (2002). Did community rating induce an adverse selection death spiral? Evidence from New York, Pennsylvania, and Connecticut. *American Economic Review* **92**, 280-294.
- Buchmueller, T., Couffinhal, A., Grignon, M. and Perronnin, M. (2004). Access to physician services: Does supplemental insurance matter? Evidence from France. *Health Economics* **13**, 669–687.
- Buzzacchi, L. and Valletti, T. M. (2005). Strategic price discrimination in compulsory insurance markets. *The Geneva Risk and Insurance Review* **30**, 71-97.
- Cardon, J.H. and Hendel, I. (2001). Asymmetric information in health insurance: Evidence from the National Health Expenditure Survey. *RAND Journal of Economics* **32**, 408–427.

- Chiappori, P.A. (2000). Econometric models of insurance under asymmetric information. In Dionne (ed.), *Handbook of Insurance*. 365-393. Boston: Kluwer Academic Publishers.
- Chiappori, P.A. (2006). The welfare effects of predictive medicine. In Chiappori, P.A. & Gollier, C. (eds.) *Competitive failures in insurance markets: Theory and policy implications*. 55-79. Cambridge, MA: MIT Press.
- Chiappori, P.A. and Salanié, B. (2000). Testing for asymmetric information in insurance markets. *Journal of Political Economy* **108**, 56–78.
- Chiappori, P.A. and Salanié, B. (2003). Testing contract theory: A survey of some recent work. In Dewatripont, M., Hansen, L.P. & Turnovsky, S. (eds.) *Advances in economics and econometrics: Theory and applications 1*, Eighth World Congress. 115-149. Cambridge: Cambridge University Press.
- Chiappori, P.A., Jullien, B., Salanié, B. and Salanié, F. (2006). Asymmetric information in insurance: General testable implications. *Rand Journal of Economics* **37**, 783–798.
- Cohen, A. (2005). Asymmetric information and learning in the automobile insurance market. *The Review of Economics and Statistics* **87**, 197–207.
- Cohen, A. and Siegelman, P. (2010). Testing for adverse selection in insurance markets. *Journal of Risk and Insurance* **77**, 39-84.
- Crocker, K.J. and Snow, A. (1985a). The efficiency of competitive equilibria in insurance markets with asymmetric information. *Journal of Public Economics* **26**, 207-219.
- Crocker, K.J. and Snow, A. (1985b). A simple tax structure for competitive equilibrium and redistribution in insurance markets with asymmetric information. *Southern Economic Journal* **51**, 1142-1150.
- Crocker, K.J. and Snow, A. (1986). The Efficiency Effects of Categorical Discrimination in the Insurance Industry. *Journal of Political Economy* **94**, 321-344.
- Crocker K. and Snow, A. (1992). The social value of hidden information in adverse selection economies. *Journal of Public Economics* **48**, 317-347.
- Crocker, K.J. and Snow, A. (2000). The theory of risk classification. In Dionne (ed.), *Handbook of Insurance*. 245-276. Boston: Kluwer Academic Publishers.
- Cutler, D., Finkelstein, A., and McGarry, K. (2008). Preference heterogeneity and insurance markets: explaining a puzzle. *American Economic Review Papers and Proceedings* **98**, 157-162.
- Cutler, D.M. and Reber, S.J. (1998). Paying for health insurance: the trade-off between competition and adverse selection. *The Quarterly Journal of Economics*, **113**, 433-466.
- Cutler, D.M. and Zeckhauser, R.J. (1998). Adverse selection in health insurance. *Forum for Health Economics & Policy* **1**: Article 2.

- Cutler, D.M. and Zeckhauser, R.J. (2000). The anatomy of health insurance. In Culyer, A.J. & Newhouse, J.P. (eds.) *Handbook of health economics* **1**, 563–643. Amsterdam: Elsevier Science.
- Dionne, G. and Doherty, N.A. (1994). Adverse selection, commitment and renegotiation: Extension to and evidence from insurance markets. *Journal of Political Economy* **102**, 210–235.
- Dionne, G., Doherty, N.A. and Fombaron, N. (2000). Adverse selection in insurance markets. In Dionne, G. (ed.), *Handbook of insurance*. 185–243. Boston: Kluwer Academic Publishers.
- Dionne, G. and Eeckhoudt, L. (1985). Self-insurance, self-protection and increased risk aversion. *Economics Letters*, 39-43.
- Dionne, G. and Gagné, R. (2001). Deductible contracts against fraudulent claims: Evidence from automobile insurance. *The Review of Economics and Statistics* **83**, 290-301.
- Dionne, G., Gouriéroux, C. and Vanasse, C. (2006). Informational Content of Household Decisions with Applications to Insurance under Asymmetric Information. In Chiappori, P.A. & Gollier, C. (eds.) *Competitive failures in insurance markets: Theory and policy implications*. 159-184. Cambridge, MA: MIT Press.
- Dionne, G., Gouriéroux, C. and Vanasse, C. (2001). Testing for evidence of adverse selection in the automobile insurance market: A comment. *Journal of Political Economy* **109**, 444–473.
- Dionne, G., Maurice, M., Pinquet, J. and Vanasse, C. (2011). Incentive mechanisms for safe driving: A comparative analysis with dynamic data. *The Review of Economics and Statistics* **93**, 1, 218-227.
- Dionne, G., Michaud, P.C. and Dahchour, M. (2010). Separating moral hazard from adverse selection and learning in automobile insurance: Longitudinal evidence from France. Working paper 10-05. Canada Research Chair in Risk Management, HEC Montreal, <http://neumann.hec.ca/gestiondesrisques/10-05.pdf>. *Journal of the European Economic Association* (forthcoming).
- Dionne, G., St-Amour, P. and Vencatachellum, D. (2009). Asymmetric information and adverse selection in Mauritian slave auctions. *Review of Economic Studies* **76**, 1269-1295.
- Doherty, N. A. and Posey, L.L. (1998). On the value of a checkup: Adverse selection, moral hazard and the value of information. *Journal of Risk and Insurance* **65**, 189-211.
- Doherty, N.A. and Thistle, P.D. (1996). Adverse selection with endogenous information in insurance markets. *Journal of Public Economics* **63**, 83-102.
- Dubey, P. and Geanakoplos, J. (2002). Competitive pooling: Rothschild-Stiglitz reconsidered. *The Quarterly Journal of Economics* **117**, 1529-1570.
- Einav, L. and Finkelstein, A. (2011). Selection in Insurance markets: Theory and empirics in pictures. *Journal of Economic Perspectives* **26**, 1-17.

- Ettner, S.L. (1997). Adverse selection and the purchase of medigap insurance by the elderly. *Journal of Health Economics* **16**, 543–562.
- Fang, H., Keane, M.P. and Silverman, D. (2008). Sources of advantageous selection: Evidence from the medigap insurance market. *Journal of Political Economy* **116**, 303–350.
- Finkelstein, A., and McGarry, K. (2006). Multiple dimensions of private information: Evidence from the long-term care insurance market. *American Economic Review* **96**, 938–958.
- Finkelstein, A. and Poterba, J. (2006). Testing for Asymmetric Information Using “Unused Observables.” NBER Working Paper 12112.
- Finkelstein, A., Poterba, J. and Rothschild, C. (2009). Redistribution by insurance market regulation: Analyzing a ban on gender-based retirement annuities. *Journal of Financial Economics* **91**, 38-58.
- Gouriéroux, C. (1999). The econometrics of risk classification in insurance. *The Geneva Papers on Risk and Insurance Theory* **24**, 2, 119-137.
- Harrington, S. (2010a). The health insurance reform debate. *Journal of Risk and Insurance* **77**, 1, 5-38.
- Harrington, S. (2010b). U.S. health-care reform: The patient protection and affordable care act. *Journal of Risk and Insurance* **77**, 3, 703-708.
- Harsanyi, J.C. (1953). Cardinal utility in welfare economics and in the theory of risk-taking. *Journal of Political Economy* **61**, 434-435.
- Harsanyi, J.C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy* **63**, 309-321.
- Hellwig, M.F., 1987. Some recent developments in the theory of competition in markets with adverse selection. *European Economic Review* **31**, 319-325.
- Hendel, I. and Lizzeri, A. (2003). The role of commitment in dynamic contracts: Evidence from life insurance. *Quarterly Journal of Economics* **118**, 299–327.
- Hoel, M., Iversen, T., Nilssen, T. and Vislie, J. (2006) Genetic testing in competitive insurance markets with repulsion from chance: A welfare analysis. *Journal of Health Economics* **26**, 251-262.
- Hoel, M. and Iversen, T. (2000). Genetic testing when there is a mix of compulsory and voluntary health insurance. *Journal of Health Economics* **21**, 253-270.
- Holmstrom, B. (1979). Moral hazard and observability. *Bell Journal of Economics* **10**, 74-91.
- Hoy, M. (1982). Categorizing risks in the insurance industry. *Quarterly Journal of Economics* **97**, 321-336.
- Hoy, M. (2006). Risk classification and social welfare. *The Geneva Papers* **31**, 245–269.

- Hoy, M. and Polborn, M. (2000). The value of genetic information in the life insurance market. *Journal of Public Economics* **78**, 235-252.
- Hoy, M. and Witt, J. (2007). Welfare effects of banning genetic information in the life insurance market: The case of BRCA1/2 Genes. *Journal of Risk and Insurance* **74**, 523–546.
- Martin, A. (2007). On Rothschild-Stiglitz as competitive pooling. *Economic Theory* **31**, 371-386.
- Miyazaki, H. (1977). The rat race and internal labor markets. *Bell Journal of Economics* **8**, 394-418.
- Netzer, N. and Scheuer, F. (2008). Competitive markets without commitment. Working Papers 0814, University of Zurich, Socioeconomic Institute.
- Pauly, M.V. (1974). Overinsurance and public provision of insurance: The roles of moral hazard and adverse selection. *Quarterly Journal of Economics* **88**, 44-62.
- Pauly, M.V., Withers, K.H., Viswanathan, K.S., Lemaire, J., Hershey, J.C., Armstrong, K. and Asch, D.A. (2003). Price elasticity of demand for term life insurance and adverse selection, NBER Working paper no 9925.
- Picard, P. (2009). Participating insurance contracts and the Rothschild-Stiglitz equilibrium puzzle. Working Papers hal-00413825, HAL.
- Puelz, R. and Snow, A. (1994). Evidence on adverse selection: Equilibrium signaling and cross-subsidization in the insurance market. *Journal of Political Economy* **102**, 236–257.
- Richaudeau, D. (1999). Automobile insurance contracts and risk of accident: An empirical test using French individual data. *The Geneva Papers on Risk and Insurance Theory* **24**, 97–114.
- Riley, J.G., (1979). Informational Equilibrium. *Econometrica* **47**, 331-359.
- Rothschild, C. (2011). The efficiency of categorical discrimination in insurance markets. *Journal of Risk and Insurance*, **78**, 267–285.
- Rothschild, M and Stiglitz, J. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics* **90**, 629-649.
- Shavell, S. (1979). On moral hazard in insurance. *The Quarterly Journal of Economics* **93**, 541-562.
- Spence, M. (1978). Product differentiation and performance in insurance markets. *Journal of Public Economics* **10**, 427-447.
- Stiglitz, J. (1977). Monopoly, non-linear pricing and imperfect information: The insurance market. *Review of Economic Studies* **44**, 407-430.
- Strohmenger R. and Wambach, A. (2000). Adverse selection and categorical discrimination in the health insurance market: the effects of genetic tests. *Journal of Health Economics* **19**, 197-218.

Villeneuve, B. (2003). Mandatory pensions and the intensity of adverse selection in life insurance markets. *The Journal of Risk and Insurance* **70**, 527-548.

Villeneuve, B. (2005). Competition between insurers with superior information. *European Economic Review* **49**, 321-340.

Winter, R. (2000). Optimal insurance under moral hazard. In Dionne (ed.), *Handbook of Insurance*. 155-183. Boston: Kluwer Academic Publishers.

Wilson, C. (1977). A model of insurance markets with incomplete information. *Journal of Economic Theory* **16**, 167-207.

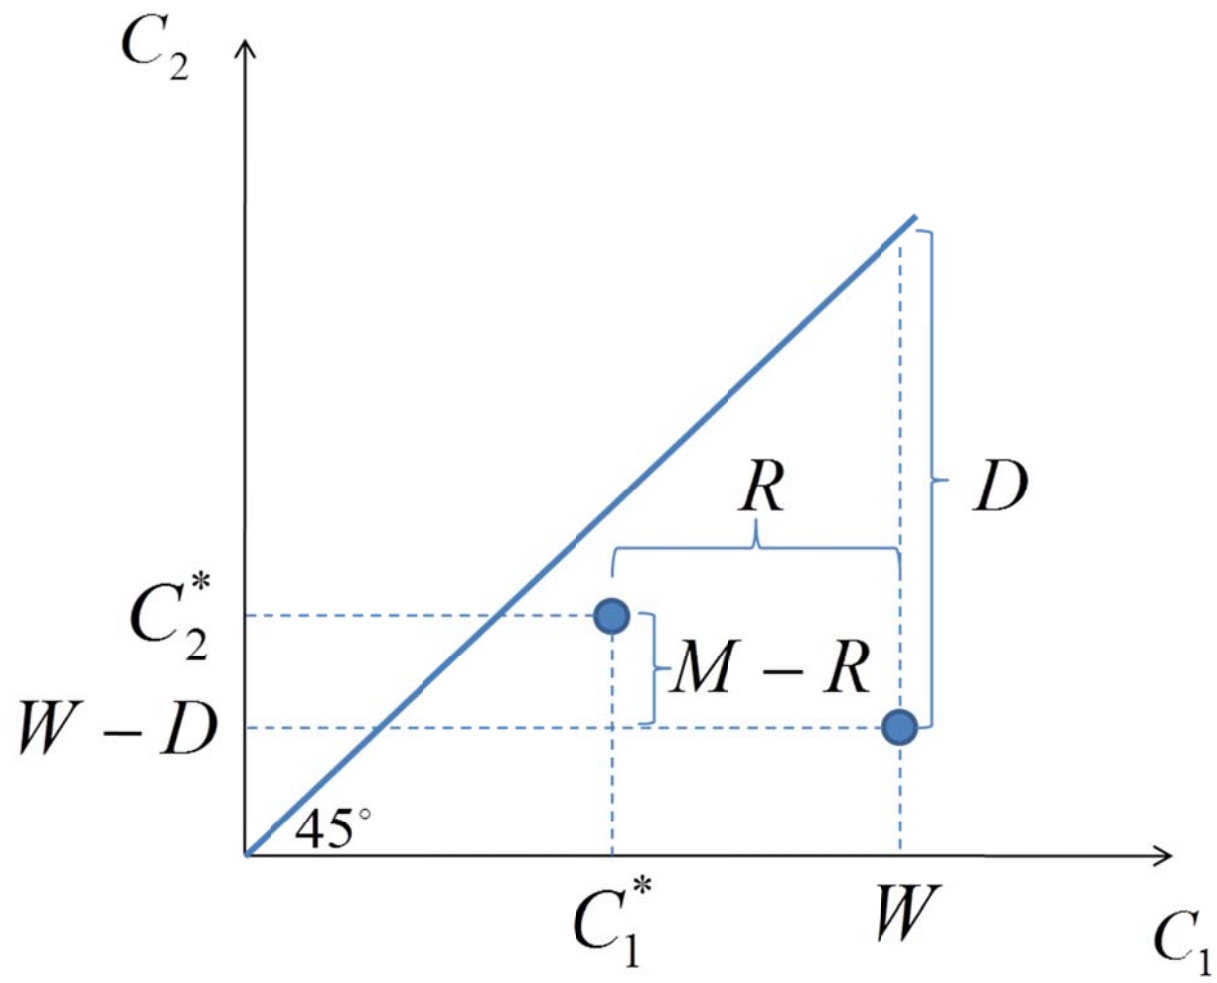


Figure 1: Basic Framework

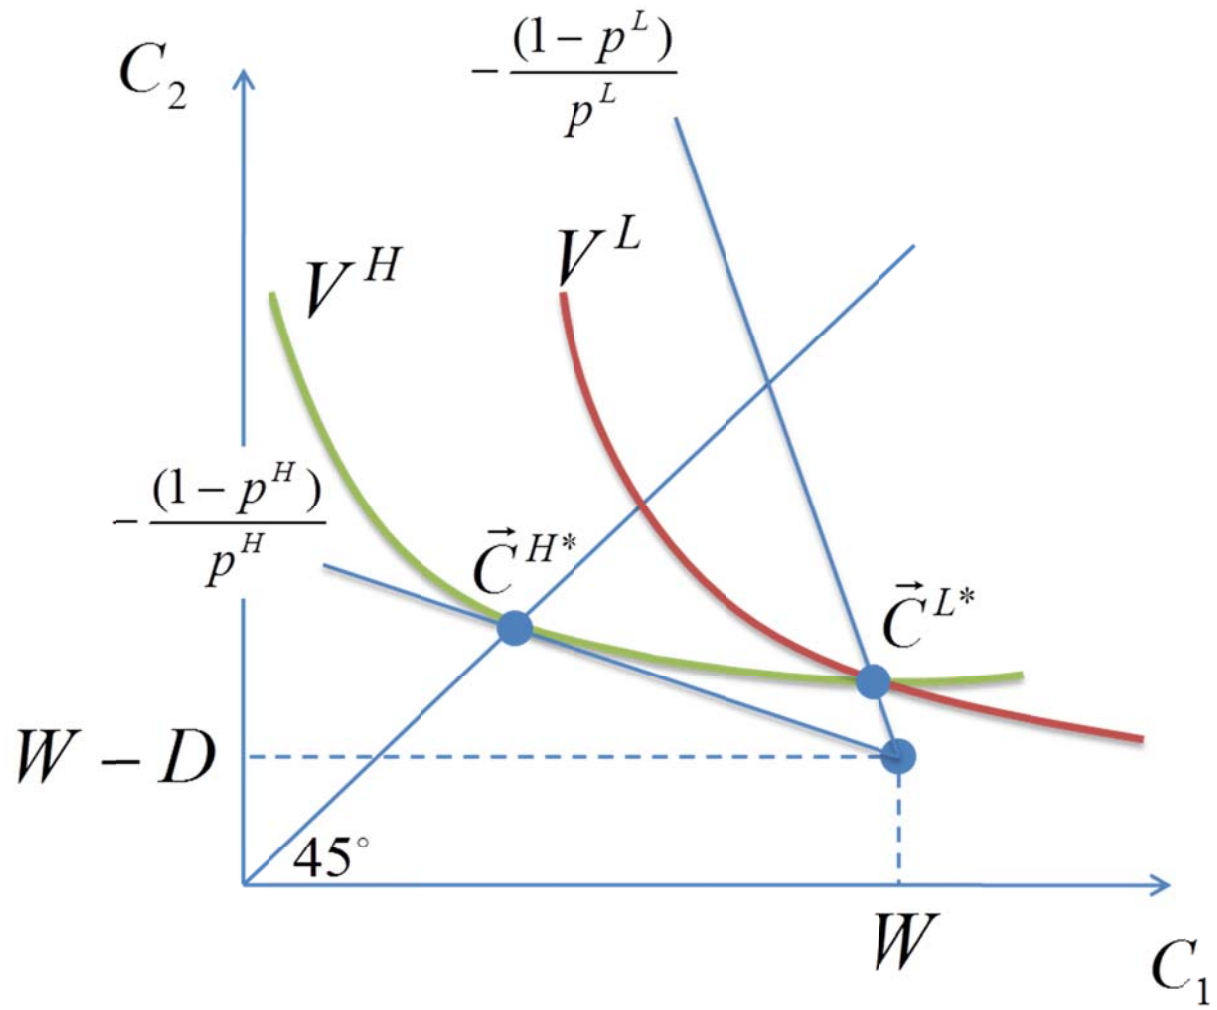


Figure 2: Rothschild-Stiglitz Equilibrium Candidate

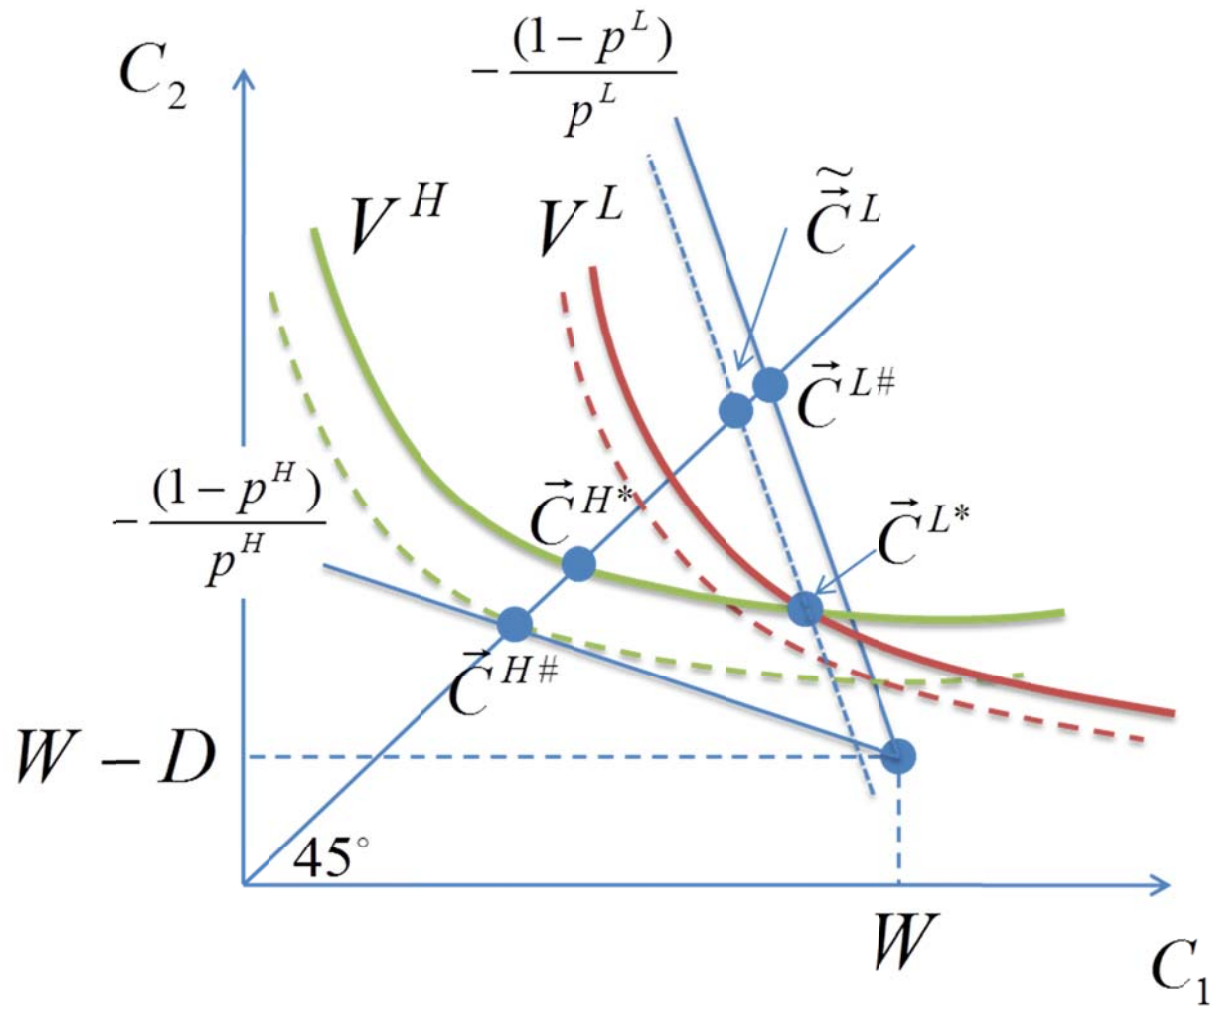


Figure 3: Miyazaki-Wilson-Spence Equilibrium with Cross-Subsidies

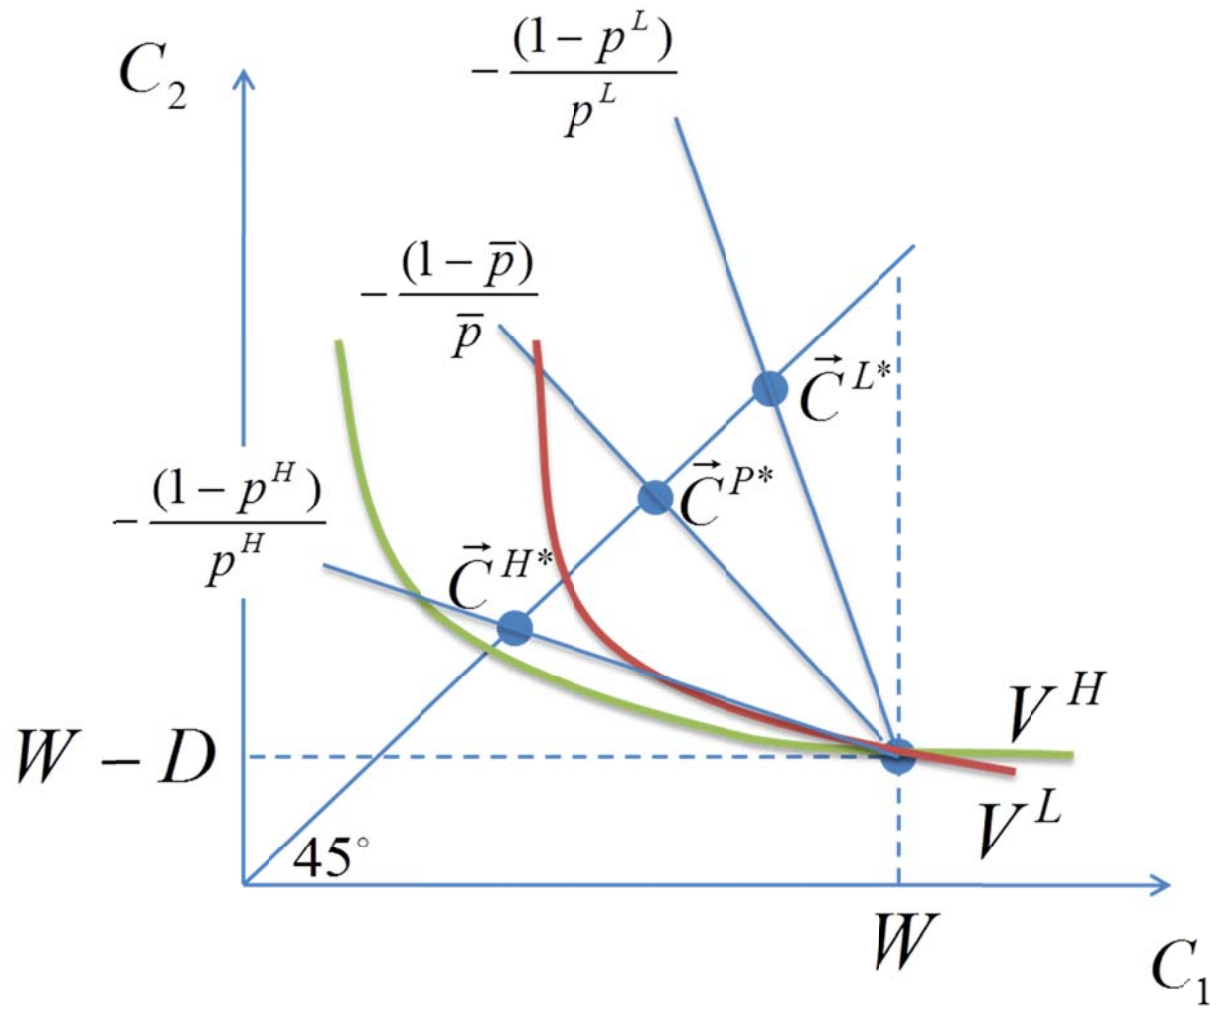


Figure 4: Risk Classification can have Purely Distributional Consequences

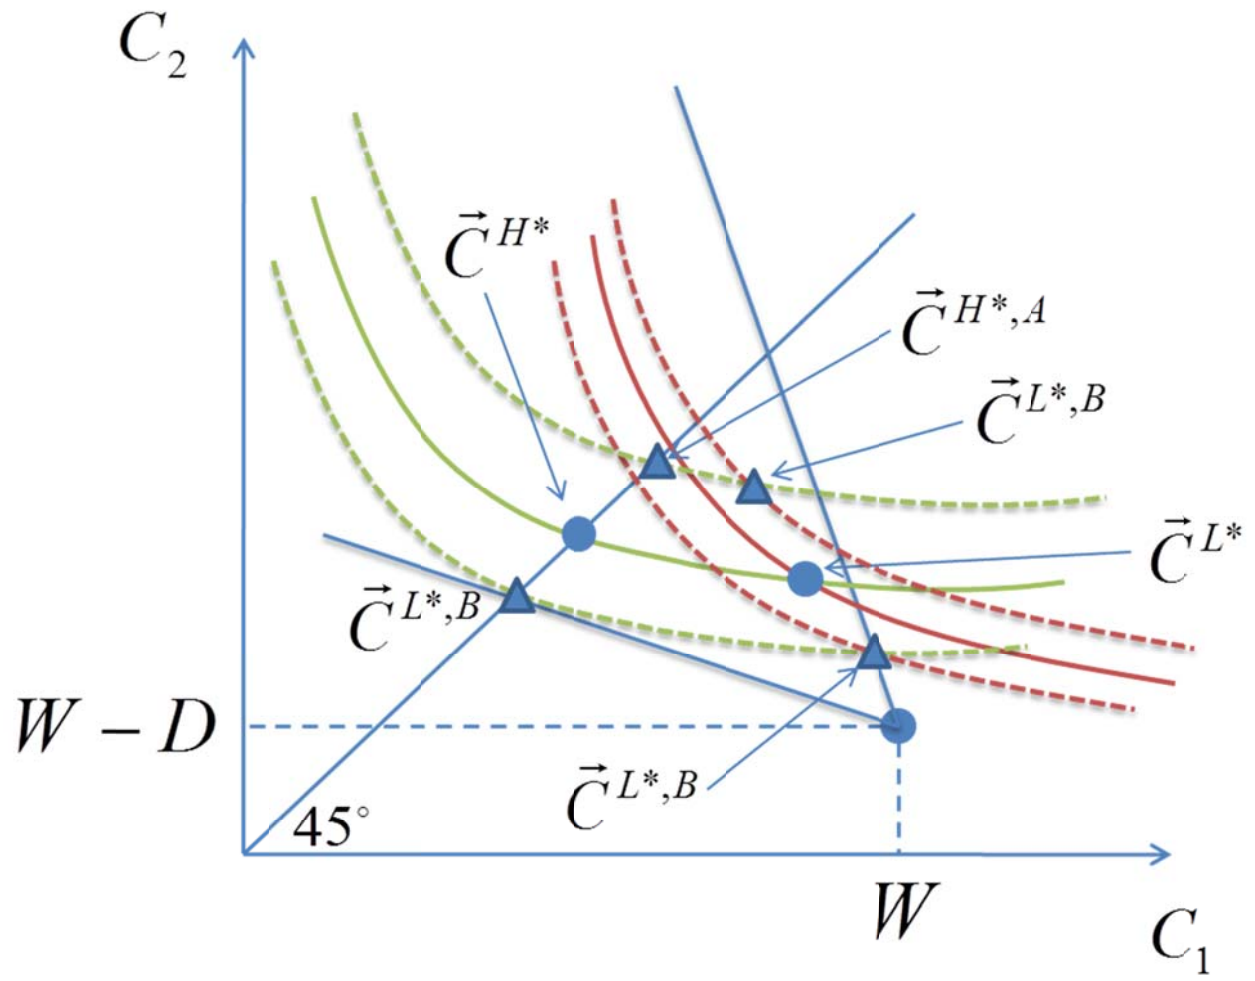


Figure 5: Risk Classification can have both Efficiency and Distributional Consequences

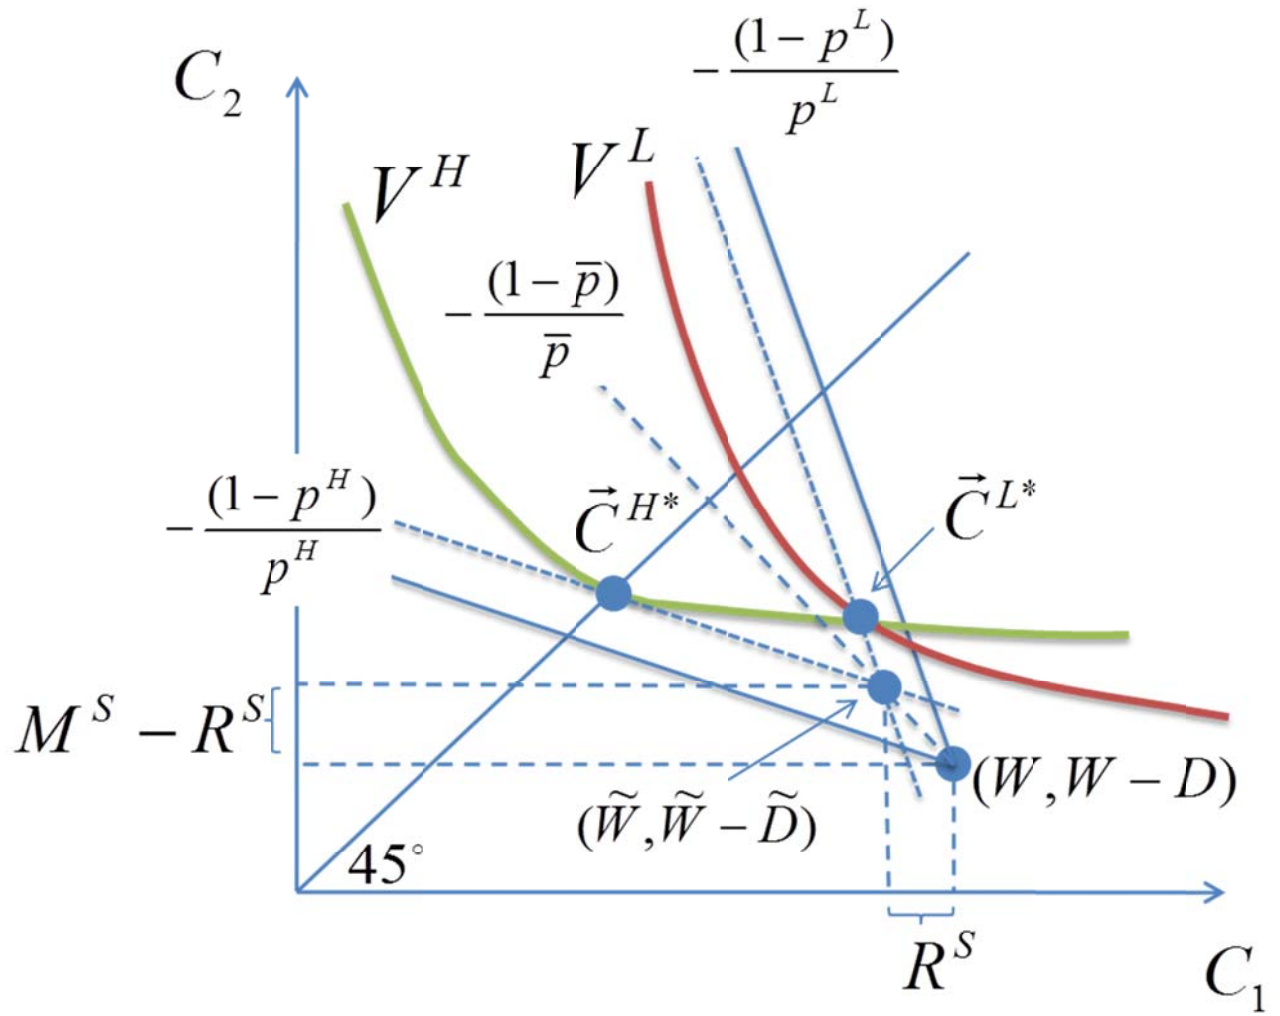


Figure 6: Partial Social Insurance Provision is Preferable to a Ban on Risk Classification

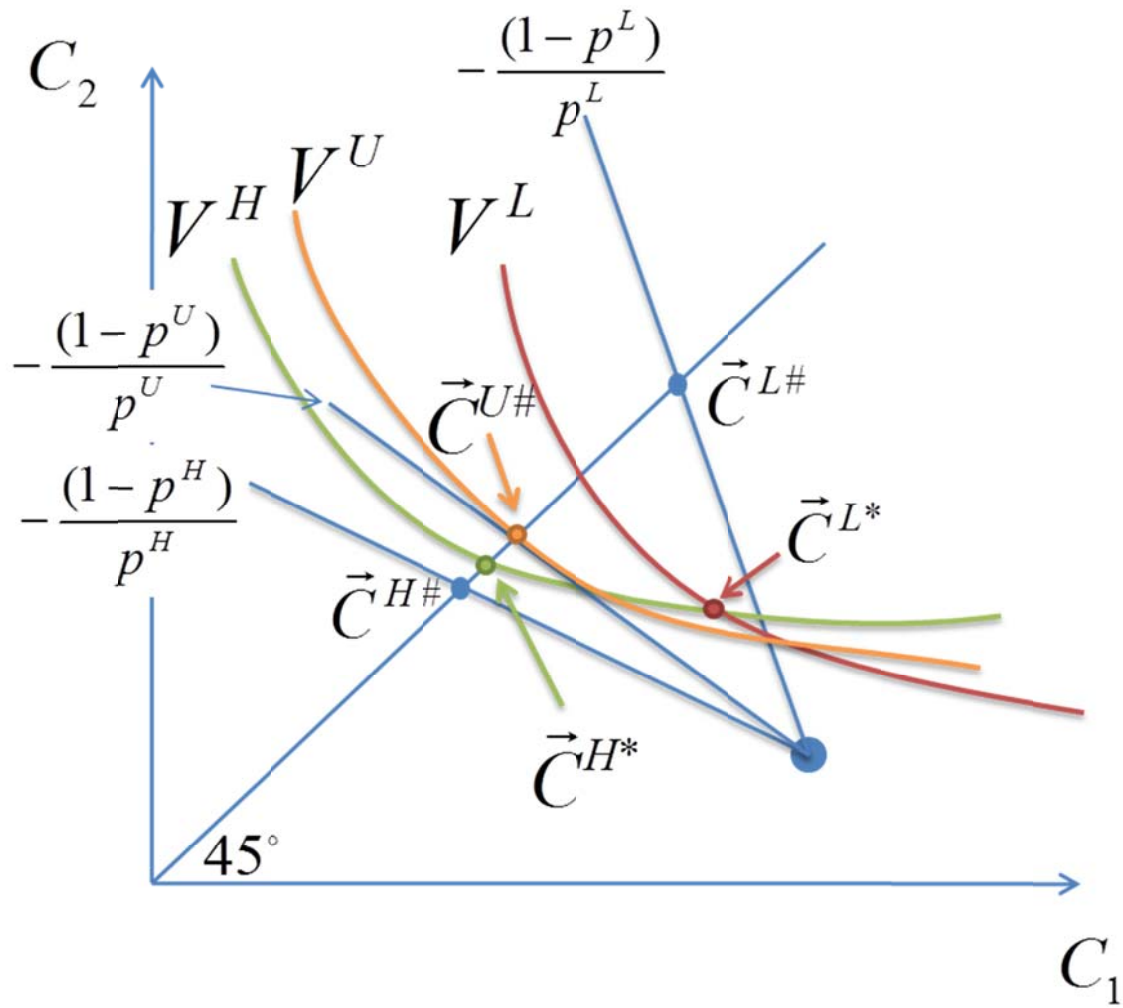


Figure 7: Market Outcomes with Private Test Results and Public Information Status

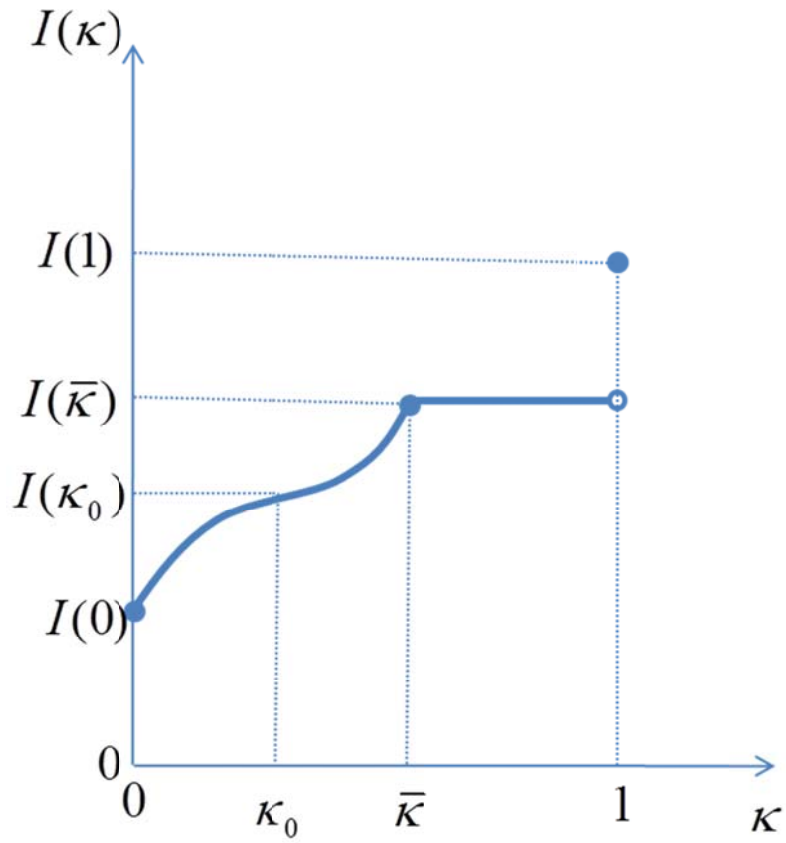


Figure 8: The Value of Information when Tests are Unobservable but Verifiable

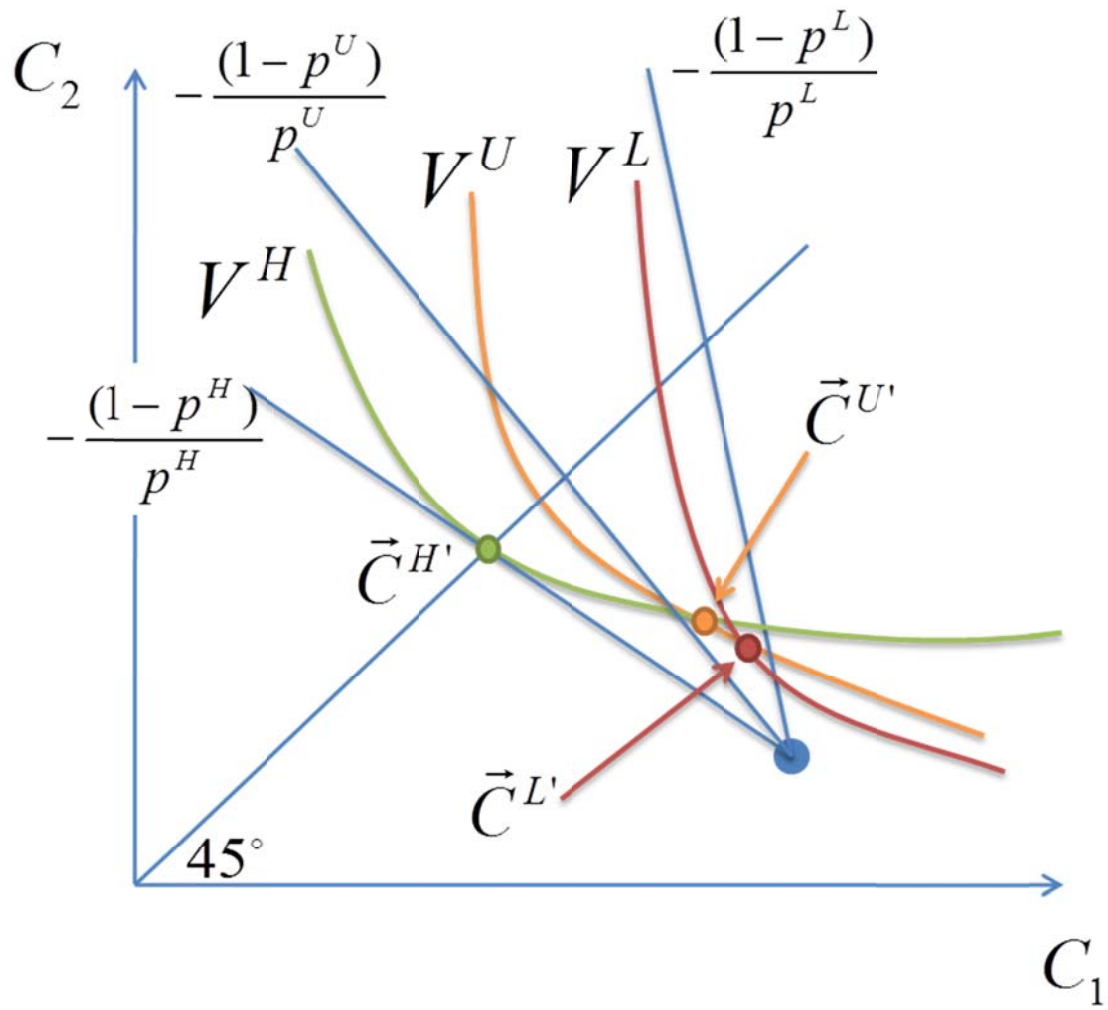


Figure 9: An MWS Market Outcome with Three Types

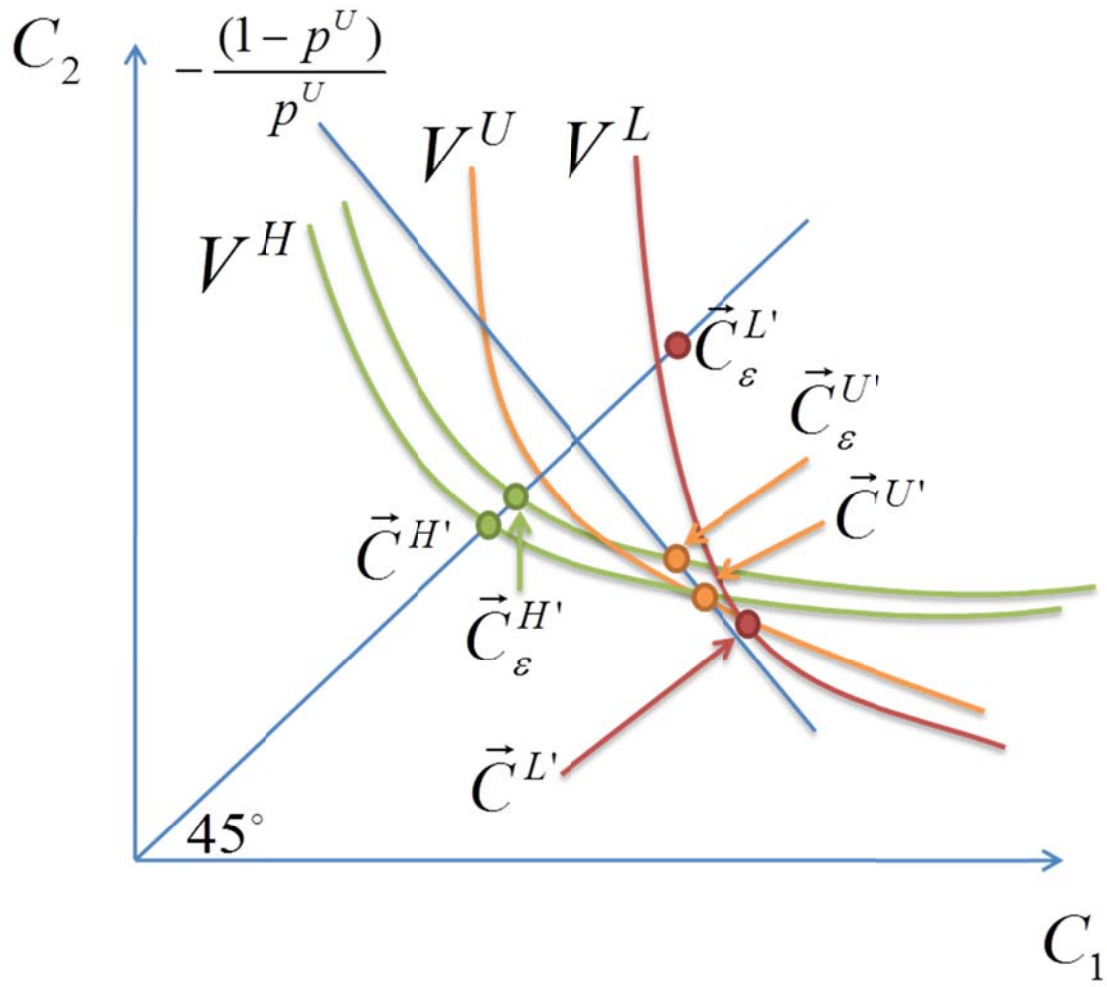


Figure 10: The Inefficiency of Banning Risk Classification with Endogenous Information

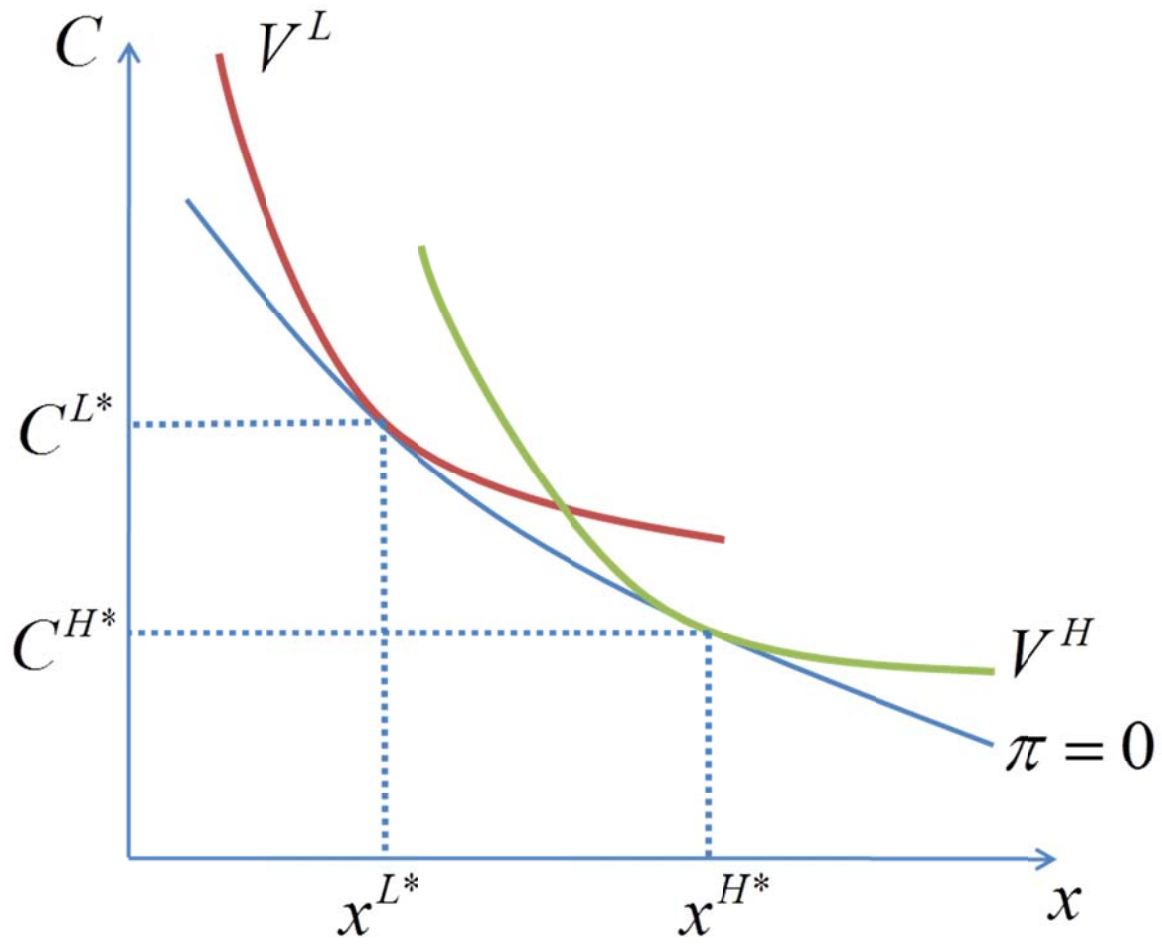


Figure 11: First-Best Break-Even Allocations with Endogenous Risk Classification