

Some Novel Perspectives on Risk Classification

R. Guy Thomas

Institute of Mathematics, Statistics & Actuarial Science, University of Kent, Canterbury CT2 7NF, U.K.
E-mail: r.g.thomas@kent.ac.uk

This paper considers a number of novel perspectives on risk classification, primarily in the context of life and critical illness insurance. I suggest that the terminology of “adverse selection” is often misleading, because from a public policy viewpoint, adverse selection may not always be adverse. I suggest that public policymakers should consider the criterion of “loss coverage”, and that in many markets a socially optimal level of adverse selection is that which maximises loss coverage. A review of empirical studies suggests that adverse selection is often difficult to observe in practice; this leads to the concept of propitious selection, and various psychological perspectives on risk classification. I suggest that competition between insurers in risk classification can sometimes be characterised as a malevolent invisible hand, and that public policy should direct competition towards areas that are more clearly beneficial to all insurance customers. I also consider the perspectives of risk classification as blame, the conflict between risk classification and human rights, and the fallacy of the one-shot gambler.

The Geneva Papers (2007) 32, 105–132. doi:10.1057/palgrave.gpp.2510118

Keywords: adverse selection; propitious selection; loss coverage; competition; psychology; one-shot gambler

Introduction

This paper considers a number of novel – or at least, seldom-aired – perspectives on risk classification. These perspectives arise mainly by considering risk classification and its consequences from viewpoints other than those typical of insurance actuaries or economists, who have tended (as others have recently noted¹) always to take the insurer’s point of view.

Much research and discussion on risk classification and public policy is affected by a structural bias, in that it tends to be funded or influenced by the insurance industry. Money not only buys political influence; used cleverly, it can also buy intellectual influence. This is not always because the insurance industry exercises direct control over research and publication (although in many cases it does). It is also because the career-minded researcher recognises, implicitly or explicitly, that the surest route to research grants, publication, promotion and peer esteem is to conform to an industry perspective – or at least, not to pursue an overtly critical one. Through circumstances rather than virtue, I am indifferent to these considerations (although no doubt prone to others).

¹ For example Baker (2003).

Some of the ideas in this paper were originally developed in a response² to a public consultation³ on genetics and risk classification in life and critical illness insurance⁴ initiated by the Human Genetics Commission in the United Kingdom. In this context, I have largely ignored issues of possible moral hazard, which may well loom larger in other classes of insurance. I have also largely ignored the possibility of government transfer payments towards higher risks as a remedy for perceived inequities arising from risk classification.⁵ This latter assumption is justified partly by the fact that government transfers related to life insurance, whether in the form of premium vouchers or direct provision of benefits, seem in most jurisdictions to be non-existent or meagre. Another justification is that, at least in the policy world, organisations (such as many insurance companies) which argue that risk classification by insurers should be unregulated also seek, in a broader context, to argue that private insurance is superior to social insurance and that the State's role in direct provision of benefits should be restricted or reduced. Viewed in this broader context, the suggestion of government transfers as a remedy for inequities arising from risk classification often appears ineffectual or insincere.

The remainder of the paper is presented as follows. The next section discusses the idea that adverse selection may not be adverse, and suggests that public policymakers concerned with risk classification should consider the criterion of “loss coverage”; the first third of the paper explains and develops this concept. I then briefly review some empirical studies of adverse selection, which suggest that the phenomenon is often rather difficult to observe in practice; this leads me to consider propitious selection, and various psychological perspectives on risk classification. I go on to suggest that competition between insurers in risk classification can sometimes be characterised as a malevolent invisible hand, and that perhaps public policy should direct competition towards areas that are more clearly beneficial to all insurance customers. The final third of the paper considers three further perspectives: risk classification as blame; the conflict between risk classification and human rights; and the fallacy of the one-shot gambler. The last section summarises the paper.

Optimal (adverse) selection: maximising loss coverage

Adverse to whom?

Much discussion of risk classification is concerned with the phenomenon of *adverse selection*. Almost invariably, the discussion is framed negatively: adverse selection is conceived as a phenomenon that should be deprecated, avoided or minimised. This framing and vocabulary reflect a structural bias: “adverse” selection primarily and most obviously means *adverse to an insurer*.

² Thomas (2001).

³ Human Genetics Commission (2001).

⁴ Critical illness insurance is a common design of long-term policy in the United Kingdom, which pays a fixed cash sum on diagnosis of any of a list of illnesses (heart attack, cancer, stroke, multiple sclerosis, etc.).

⁵ For example, as considered in Crocker and Snow (1986).

From a public policy viewpoint, it is not obvious that adverse selection is always adverse. From a public policy viewpoint, some types of insurance are probably social (or merit) goods,⁶ which public policy should ensure is available widely. A degree of adverse selection in some insurance markets may then be seen as desirable: it means that the right people, people more likely to suffer loss, tend to buy insurance.

The extent to which adverse selection is seen as desirable from a public policy viewpoint may depend on the insurance market in question. For example, some people might regard a degree of adverse selection in life or health insurance as positive; but the same people might (or might not) regard adverse selection in motor insurance as negative. However, the important point is that the assumptions implicit in much commentary – that adverse selection is always unwelcome, and that policy should be directed towards its elimination – may be inappropriate from a public policy viewpoint.

The economics literature often suggests that adverse selection is also adverse in a different sense, in that it reduces the efficiency of contracting. In the classical model of insurance under asymmetric information, it leads higher risks to receive full coverage at an actuarially fair price, while lower risks receive less than full coverage, so that a Pareto improvement would be possible from accurate risk classification.⁷ A similar result obtains under some alternative concepts of equilibrium.⁸ But this type of theorising often seems at best tenuously connected with the reality of life insurance markets, in which higher risks are charged higher prices or excluded from insurance, and there is no rationing of insurance for lower risks.

The economics literature also tends to discuss alleged welfare losses from adverse selection in terms that can seem rather abstract and theoretical, with subtly different answers depending on which particular concept of game-theoretic equilibrium is selected; in contrast, the welfare losses of people excluded from insurance are often human, compelling and observable. This does not necessarily mean that the abstractions should be dismissed; but it does mean that one should consider carefully to what extent the abstractions represent real markets and real human phenomena.

Maximising loss coverage

If some degree of adverse selection may be desirable in certain markets, the question arises of how much? I suggest that in many markets, a public policymaker should seek to maximise the “overlap” of insurance coverages and loss events – the “loss coverage” – that is, to maximise the proportion of loss events that is covered by insurance.⁹

⁶ That is, goods the supply of which generates positive externalities. For example, life insurance may reduce the need for government welfare transfers to bereaved dependants, facilitate entrepreneurial risk-taking, etc.

⁷ For example Rothschild and Stiglitz (1976).

⁸ For example Wilson (1977).

⁹ As noted in the introduction to the paper, I assume that moral hazard is small enough to be ignored. I also assume that the restitution of losses through insurance is regarded as social good, which public policy should seek to promote. These assumptions seem appropriate for life insurance. Both may be less appropriate for certain other types of insurance.

One way of maximising loss coverage is to make insurance compulsory. This approach is sometimes taken, particularly where the purpose of the insurance is to protect the interests of unidentified third parties or the public at large; in other words, the insurance is largely a social good rather than a private good. Compulsory employers' liability insurance and third-party motor insurance are examples of this in the U.K. In these cases, the perceived societal need for and benefit from the insurance are little affected by circumstances or preferences of the insured. The compulsory approach is therefore seen to be appropriate.

But there are few types of insurance that are almost entirely social goods. Most types of insurance, such as life insurance, are seen as providing some benefit to wider society; that is, they are seen *partly* as social goods. However, the social good is achieved by means of providing *private* benefit to the insured or their immediate family. In these cases, the need for and benefit from the insurance are directly affected by personal circumstances and preference of the insured. For example, a compulsory requirement to purchase life insurance would be of no benefit to single persons with no dependants. In these cases, the compulsory approach may be seen as an undesirable restriction of individual choice.

Although insurances that provide a mixture of private and social good are usually not seen as appropriate cases for compulsion, this does not mean that a policymaker should take no interest whatsoever in their provision. If it is accepted that the insurance is at least *partly* a social good, the policymaker may wish to promote its availability, and, in particular, its availability to persons who are more likely to suffer loss. The policymaker might wish to target adverse selection high enough that many higher risks can obtain coverage, but not so high that an adverse selection spiral develops. This may mean targeting a higher level of adverse selection than insurers might wish. Informally, we can say that the public policymaker wants "*insurance bought by people likely to need it*".

There may be a few insurance markets where insurance is *not* seen as having any social good element, or where the social good element is seen as negligible. A clear example would be insurance against the cost of legal or regulatory penalties; a less-clear example might be insurance of pet animals. In these cases, a policymaker would probably place a much lower value on maximising loss coverage than in markets such as life insurance. Indeed, the policymaker might even seek to *minimise* loss coverage, by discouraging the insurance. In extreme cases, this might mean banning the insurance.¹⁰

A simple example

The concept of adjusting the degree of risk classification to maximise loss coverage is illustrated in the following model. There are two risk groups: a lower risk group with risk 1/100 and a higher risk group with risk 4/100. All insurance contracts are for a fixed sum assured.¹¹ In Scenario 1 (shown as Table 1), actuarially fair premiums are

¹⁰ For example, the Financial Services Authority (FSA) in the U.K. prohibits regulated firms from insuring against the cost of FSA penalties.

¹¹ This allows adverse selection to be represented by varying proportional take-up of insurance in each of the two groups. The assumption simplifies the presentation, but it is not necessary.

charged to each risk group. In Scenario 2 (shown as Table 2), insurers are not allowed to discriminate; there is a single premium rate irrespective of risk group, and some adverse selection occurs. Take-up of insurance increases in the higher risk group and reduces in the lower risk group, and a slightly lower total number of policies are issued. The net effect is favourable, in that a larger proportion of the loss events is covered by insurance, that is, the loss coverage is increased. In Scenario 3 (shown as Table 3), insurers are again not allowed to discriminate, but the adverse selection is more severe: coverage shifts further towards the higher risk group, the number of policies reduces more drastically. The net effect is unfavourable, in that the loss coverage is reduced.

In my view a policymaker concerned with public welfare should regard Scenario 2, with higher overlap of insurance coverages and loss events, as superior to Scenario 1. That is, of these two scenarios, the case with higher adverse selection is superior. The more accurate risk classification in Scenario 1 produces a lower loss coverage.

Table 1 Scenario 1: with actuarially fair premium rates

<i>Risk group</i>	<i>Lower risk group</i>	<i>Higher risk group</i>
Population	10,000	2,000
Risk	1/100	4/100
Premium required	1/100	4/100
Take-up (members of group purchasing insurance)	5,000	500
(A) Expected loss events	100	80
(B) Loss events covered	50	20
Loss coverage: $\Sigma B/\Sigma A$	0.39	

25 per cent take-up in higher risk group (insurance is prohibitively expensive).
50 per cent take-up of insurance in lower risk group.
5,500 policies issued.

Table 2 Scenario 2: single premium rate

<i>Risk group</i>	<i>Lower risk group</i>	<i>Higher risk group</i>
Population	10,000	2,000
Risk	1/100	4/100
Premium required	1.60/100	
Take-up (members of group purchasing insurance)		
NB some adverse selection	4,000	1,000
(A) Expected loss events	100	80
(B) Loss events covered	40	40
Loss coverage: $\Sigma B/\Sigma A$	0.44	

Some adverse selection.
Take-up increased to 50 per cent in higher risk group.
Take-up fallen to 40 per cent take-up in lower risk group.
Reduction to 5,000 in number of policies issued.
Higher loss coverage.

Table 3 Scenario 3: single premium rate

<i>Risk group</i>	<i>Lower risk group</i>	<i>Higher risk group</i>
Population	10,000	2,000
Risk	1/100	4/100
Premium required	1.86/100	
Take-up (members of group purchasing insurance)		
NB extreme adverse selection, greatly reduced total coverage	2,500	1,000
(A) Expected loss events	100	80
(B) Loss events covered	25	40
Loss coverage: $\Sigma B/\Sigma A$	0.36	

More adverse selection.

50 per cent take-up in higher risk group (same as Scenario 2).

25 per cent take-up in lower risk group (cf. 40% in Scenario 2).

Reduction to 3,500 in number of policies issued (cf. 5,000 in Scenario 2).

Lower loss coverage than either of Scenarios 1 or 2.

Note that although slightly fewer policies are issued in Scenario 2 than in Scenario 1, *more policies are bought by the right people*, that is, people more likely to suffer loss. The loss coverage criterion captures the net effect of the shift in coverage towards higher risks (good) and the reduction in number of policies issued (bad).

On the loss coverage criterion, Scenario 3 (see Table 3) is inferior to both Scenario 1 and Scenario 2. Although the take-up of insurance is higher in the higher risk group (good), this is more than offset by the reduction in take-up in the lower risk group (bad). The net result is that the loss coverage is reduced.

Loss coverage: comments, caveats and limitations

It might be suggested that the lower number of policies issued in Scenario 2 as compared with Scenario 1 is a disadvantage, reflecting a lower efficiency of contracting. My view is that loss coverage is a more important public policy objective in life insurance than efficiency of contracting.

Nothing in the discussion above has suggested that the loss coverage will *necessarily* be increased by restricting risk classification. In other words, nothing has been said about which of Scenarios 2 or 3 is more likely if restrictions on risk classification are imposed in a market where none presently exist. It is possible that in some markets and in respect of some risk factors, Scenario 3 represents the *status quo*: in other words, that a higher loss coverage could be achieved if insurers were required to classify risk more accurately than market competition leads them to do. However, competition between insurers induces a tendency to more accurate risk classification, even where this has no public benefit (this is elaborated later in the paper). This makes the idea that market competition could lead to insufficient levels of risk classification seem less plausible.

The essential difference between Scenario 2 and Scenario 3 is one of responsiveness of the lower and higher risk groups to the change in risk classification – loosely speaking, price elasticities of demand in higher and lower risk groups.¹² When restrictions on risk classification are imposed (Scenario 2), a better outcome occurs if the incremental number of higher risks that gain coverage is at least one quarter the number of lower risks that desert the market.

In our scenarios, the sizes of the lower and higher risk groups, and their respective risks, are fixed. More generally, the effect on the loss coverage of restricting risk classification will also depend on (a) the *relative* sizes of the lower and higher risk groups and (b) the *relative* risks of the lower and higher risk groups.

The above observations are intuitive, but we can formalise the presentation as follows. Define P_1 and P_2 as the lower and higher risk populations, θ_1 and θ_2 as the respective proportions of each population taking up insurance, and μ_1 and μ_2 as their respective risks. Then define

$$\text{loss coverage} = \frac{\mu_1\theta_1P_1 + \mu_2\theta_2P_2}{\mu_1P_1 + \mu_2P_2}$$

The suggested policymaker's objective can then be seen as maximising a weighted average of the proportional take-ups of insurance θ_1 and θ_2 in the lower and higher risk groups, the weights being the expected population losses μ_1P_1 and μ_2P_2 , respectively. Dividing through by μ_1P_1 , this is equivalent to maximising

$$\theta_1 + (\mu_2P_2/\mu_1P_1)\theta_2$$

from which it can be seen that the suggested policymaker's objective is parameterised by the relative risk (μ_2/μ_1) in higher and lower risk groups, and the relative size (P_2/P_1) of the groups.

In our scenarios, this objective function amounted to $\theta_1 + 0.8\theta_2$. Any restriction on risk classification is expected to lead to a reduction in θ_1 and an increase in θ_2 . The net effect is an increase in the loss coverage, provided $-\Delta\theta_1 < 0.8\Delta\theta_2$. For example:

- in Scenario 2, $-\Delta\theta_1=0.10 < 0.8\Delta\theta_2=0.2$; but
- in Scenario 3, $-\Delta\theta_1=0.25 > 0.8\Delta\theta_2=0.2$.

Note that the quantity (μ_2P_2/μ_1P_1) is the ratio of the *total expected losses* in the higher risk group relative to the lower risk group. Under the loss coverage criterion, the weight placed by the policymaker on proportional take-up in each risk group depends on the group's total expected losses.

More generally, let the population in the i th risk group be P_i , let the risk for that group be μ_i , and let the proportion of the population taking up insurance in that group

¹² I say "loosely speaking" because price elasticity of demand is usually understood as a marginal measure with respect to small changes in price. This is not readily applicable to our scenarios, where the two groups experience changes in price of different sign and over different large ranges. However, if we define elasticity of demand as $(\Delta t/t)/(\Delta p/p)$ for the high-risk group, where t and p are that group's original take-up and price in Scenario 1, and proceed analogously for the lower risk group, then the ratio of the elasticities thus defined in lower and higher risk groups is 0.20 for Scenario 2, but 0.31 for Scenario 3.

be θ_i . Then, we can define

$$\text{loss coverage} = \frac{\sum \theta_i \mu_i P_i}{\sum \mu_i P_i}$$

as a criterion for the social efficacy of insurance.¹³ The policymaker's objective is to maximise the weighted average of the take-ups θ_i , where the weights for the risk groups are the expected population losses $\mu_i P_i$.

Loss coverage in different markets

It was noted above that the change in the loss coverage if risk classification is restricted depends on the responsiveness of high- and low-risk groups to changes in the price schedule. In the extreme, if the price elasticity of demand for insurance in the lower risk group was zero, reducing risk classification would always increase the loss coverage: more higher risks would be covered, with no reduction in coverage of lower risks.

The responsiveness of higher and lower risk groups to changes in price will probably be different in different insurance markets. It follows that even if loss coverage is regarded as uniformly desirable across different markets, different levels of risk classification are socially optimal in different markets. It is difficult to say much about *which* markets: there is almost no data on price elasticities of demand for insurance, and certainly none that I know of for different elasticities in higher and lower risk groups. Price elasticity would normally be lower in the lower risk group simply because of the lower price, but this need not mean that it is *sufficiently* lower. However, note that if the higher risk group tends to be declined altogether by insurers under unrestricted risk classification, we might expect a large increase in take-up – in effect, high elasticity of (satisfied) demand – in that group when restrictions on risk classification are imposed.

Alternative public policy objectives

In the discussion above the policymaker's objective was to maximise the loss coverage, with no distinction between loss events arising from the high-risk or low-risk groups. This utilitarian perspective is a starting position, but additional or alternative policy objectives are plausible. For example, if members of the high-risk group are socially disadvantaged, the policymaker might well place a higher value on coverage of loss events in the higher risk group than in the standard risk group. The policymaker might also place value on ensuring the optional *availability* of insurance to members of the higher risk group, distinct from the actual take-up of the option. These approaches would be consistent with many other public policy initiatives directed at redressing disadvantages or social exclusion of the higher risk group, for example disability discrimination laws.

¹³ The denominator will be a constant for any given market, but is included to normalise loss coverage to the range [0,1].

Policymakers may also have other objectives that are not directly related to insurance, but instead derive from social policies or legal principles that the policymaker considers should override insurance considerations. For example, some questions that insurers wish to ask for risk classification may potentially impinge on public health, or may be regarded as unwarranted infringements of privacy. Statistical discrimination by reference to particular criteria may be deprecated in principle. Charging costs arising from disability to the disabled person may be deprecated in principle. (These ideas are discussed later in the paper.)

Nevertheless, while recognising the difficulty of practical application and that alternative or additional policy objectives are possible, the concept of loss coverage provides another way of thinking about the design of insurance systems from a public policy perspective. In many markets, loss coverage may be a useful criterion for an insurance-focused public policymaker.

Loss coverage: the insurers' perspective

The priorities of a profit-maximising insurer in relation to loss coverage might be slightly different to those of the public policymaker. Higher loss coverage (Scenario 2) is associated with higher total premium income, but the total number of policies sold is lower than under more accurate risk classification (Scenario 1). I am not sure which might be preferred by an insurer in practice.

The potential conflict between public policy objectives and an insurer's profit-maximisation objective may run deeper than this. If we relax the assumption that total premiums exactly equal total claims, we can observe that in practice an insurer might prefer to design policy conditions or selling procedures with the intention of *reducing* loss coverage. Informally, we can say that the insurer wants "*insurance sold to the people less likely to need it*", or at least less likely to make a successful claim.

The insight that insurers would rationally prefer to sell insurance to people who are unlikely to make a successful claim tends not to be articulated in the literature (structural bias). However, practical examples of the commercial success that follows from this approach can be observed. For example, payment protection insurance, which covers the repayments on a loan if the borrower becomes ill or unemployed, is often sold in the U.K. to people whose circumstances (e.g., self-employed, or contract workers) will foreseeably prevent them from being able to claim. Because the insurance is often sold to people who do not need it, it is an extremely profitable business for insurers, with profit margins around 75 per cent of premiums. Some U.K. companies are reported to have incorporated subsidiaries in offshore centres, which has the effect of shielding these profit margins from government and public scrutiny.¹⁴ Similar observations have been made regarding payment protection insurance in the U.S.¹⁵

The insight that insurers' greatest commercial successes are sometimes achieved by selling insurance to people unlikely to need it may also cast some light on the phenomenon of critical illness insurance. This is a relatively new type of policy in the

¹⁴ Inman (2004).

¹⁵ Hunt and Birnbaum (2001).

U.K: it was unknown before the 1980s, did not sell in large numbers until the late 1990s, and would perhaps not be greatly missed by consumers or policymakers if it disappeared. Insurance industry commentators often remark on the mismatch between the limited circumstances in which the policy pays out (that is, if the policyholder is diagnosed with one of a rather arbitrary list of illnesses) and policyholders' plausible real needs (that is, a payout if they suffer *any* prolonged major illness, not just an arbitrary few illnesses). However, this mismatch may be an important reason for the commercial success of the policy. At the point of sale, many customers probably perceive the policy as providing more comprehensive coverage than it really does.¹⁶ An unlucky few will find out later that many major illnesses do not result in a successful claim. Another means of limiting loss coverage in critical illness insurance, which has recently been criticised by the Scottish courts, is the practice of requesting minimal underwriting information at the point of sale and then attempting to allege trivial non-disclosure to void the policy at the point of claim.¹⁷ This mismatch between customer perceptions at the point of sale and the insurer's perception of what can be achieved through claims investigation may be an important ingredient of the policy's commercial success.¹⁸

Empirical studies of adverse selection

Is adverse selection significant?

The question “is adverse selection significant?” can cause confusion, because different groups tend to interpret it in different ways. Insurance underwriters tend to interpret it as a question about *incentives*: might a prospective insured sometimes have an incentive not to disclose adverse information about his risk status? The answer to this appears to be yes; but this framing of the question offers no insight beyond a justification for the existence of underwriters. Economists tend to interpret it as a question about *information asymmetry*: do insureds have an information advantage over the insurer, after whatever screening and sorting mechanisms the insurer chooses to apply? A third interpretation is as a question not about interactions between insurer and insured, but rather about *competition between insurers*: do innovations in risk classification by insurers often lead to the distinctive competitive dynamic of the “adverse selection spiral”?

The economists' interpretation leads to a possible test for adverse selection: investigate whether a positive relationship can be detected between risk level and amount of insurance purchased. Several recent studies along these lines have failed to

¹⁶ There are several reasons for this. For example, presentation and discussion of a list of salient conditions (“dread diseases” as they are sometimes aptly known) probably encourages the illusion of comprehensive cover. The many less salient (but not less common) conditions which might make the policyholder unable to work need not be mentioned.

¹⁷ Eassie (2006).

¹⁸ Further evidence of the mismatch between customer and insurer understanding of the policy is provided by the high percentage of critical illness claims rejected – over 20 per cent at most U.K. insurers (Dyson (2006)).

find such a relationship for life assurance in the United States,¹⁹ motor insurance in France,²⁰ and health insurance in the U.S.²¹ There appears to be some evidence for adverse selection in annuity markets in the U.K.,²² but based on the annuity features (e.g., capital protection, rate of increase) rather than the amount of the annuity. A recent paper by Siegelman provides a careful summary of about a dozen studies, which he summarises, perhaps a little boldly, as follows: “In general, the literature provides little or no support for the existence of selection due to information asymmetries in insurance markets.”²³

Another observation that casts some doubt on the “asymmetric information” variety of adverse selection stories is the linear pricing of most insurance products with respect to sum assured. If residual information asymmetry were a pervasive problem in insurance, one would expect insurers to charge higher premium rates for larger sums assured, just as market makers in financial markets quote wider spreads for larger transactions. The classical theory of adverse selection²⁴ predicts rationing of insurance for better risks. But rationing or non-linear pricing does not seem to be observed in life insurance in practice.

How can the absence of measurable information asymmetry in favour of insureds be explained? There seem to be two possible explanations. First, it could be that present levels of underwriting and risk classification are fully effective, so that after underwriting insurers have an information advantage, or at least no disadvantage, relative to their insureds. Second, it could be that insureds do retain information advantages after underwriting, but that they do not make use of this information – perhaps because they cannot, or perhaps because they have more important or interesting things to do. (Note that explanations to the effect that a researcher has not fully appreciated the subtleties of the insurer’s risk classification processes are unlikely to be valid: if a researcher can show that a subset of the insurer’s information set provides an information advantage over insureds, it follows *a fortiori* that the insurer’s full information set provides an advantage.)

The third interpretation of adverse selection is as a description of competition between insurers. In its extreme form, this is the story of the “adverse selection spiral” – a cycle of increasing prices and declining coverage that destroys an insurance market. This concept is familiar to all students of insurance; but, for such a familiar concept, evidenced examples are remarkably difficult to find. Siegelman²⁵ carefully debunks several alleged instances of adverse selection spirals as “urban legends”. There is one scenario where he considers that the “adverse selection spiral” story appears pertinent: competition between close but not identical substitutes (in Siegelman’s example, two alternative health plans with different risk classification schemes offered to members of a university). This is as economic theory would suggest: where a close substitute is

¹⁹ Cawley and Philipson (1999).

²⁰ Chiappori and Salanie (2000).

²¹ Cardon and Hendel (2001).

²² Finkelstein and Poterba (2000).

²³ Siegelman (2004, p. 1248).

²⁴ For example Rothschild and Stiglitz (1976).

²⁵ Siegelman (2004)

readily available, price elasticity of demand is likely to be high. But, in general, there are no close substitutes for life insurance.

In summary, adverse selection exists in the limited sense that some insureds have an incentive not to disclose information about their risk status. However, adverse selection in the sense of information asymmetry to the benefit of insureds after underwriting has proved difficult to detect. As a concept, adverse selection may be more useful in explaining interactions between competing insurers, rather than interactions between insurer and insureds.

Propitious selection

One of the reasons why adverse selection is difficult to observe in practice may be that purchasing (more) insurance is sometimes associated with lower risk, not higher risk. If risk aversion is sufficiently higher in lower risk insureds, the standard prediction of adverse selection theories can be reversed. This idea is described by the phrase “propitious selection”.²⁶

Examples of propitious selection

One area where the phenomenon of propitious selection appears to be well recognised (albeit not by name) is in motor insurance. It has been suggested by the Association of British Insurers that around 5 per cent of drivers in the U.K. are driving without insurance (this is illegal in the U.K.), and that these drivers are nine times more likely to be convicted of a serious driving offence. A case-control study in New Zealand, where motor insurance is *not* compulsory, found that after controlling for multiple confounding variables (e.g., age, sex, level of education, driving hours per week), uninsured drivers were around five times more likely than insured drivers to suffer a car-crash injury.²⁷ These authors also cite various descriptive surveys suggesting that uninsured drivers are more likely to engage in risky driving behaviours.

The two papers²⁸ by Hemenway outlining the concept of propitious selection give evidence for several other examples using U.S. data, including the following factors:

- for motorcyclists, a positive correlation between wearing a helmet and holding medical insurance;
- for car drivers, a positive correlation between purchase of non-compulsory liability insurance and not driving after drinking alcohol;
- for car drivers, a positive correlation between purchase of non-compulsory liability insurance and a range of health-related risk avoidance activities.

It is easy to think of ways in which propitious selection effects could occur in life and health insurance. It also seems plausible that they would usually be weaker than adverse selection effects. From a public policy viewpoint, this is fortunate – since if

²⁶ Hemenway (1990, 1992).

²⁷ Blows *et al.* (2003).

²⁸ Hemenway (1990, 1992).

propitious selection were stronger, that would mean that the wrong people would be buying (more) insurance.

Psychology of propitious selection

The economics literature usually seems to define risk aversion narrowly, to mean that the marginal utility of wealth decreases as wealth increases. Propitious selection stories require the assumption that individuals who are more risk averse according to the economist's narrow definition – and hence to the *financial* risks of injury and death – also take more care in relation to the associated physical risks. Economists often seem unwilling to contemplate assumptions of this nature, apparently for methodological reasons (such as a preference for only minimal assumptions about utility functions). But intuition suggests that the assumption might be plausible in many contexts, and there is some empirical work in psychology and econometrics – and also some insurance industry practices – which appear to support it.

For example, at least one laboratory study has reported associations between financial risk-taking, psychometric profiles and blood biochemistry.²⁹ Another study reported positive correlations between answers to a questionnaire designed to measure financial risk aversion and behaviours such as smoking, not having insurance, and choosing risky employment.³⁰ Another found that smokers select riskier jobs than non-smokers; and also that smokers in risky jobs receive less of a wage premium over smokers in non-risky jobs than the corresponding premium for non-smokers.³¹ From a more psychological perspective, it has been suggested³² that many decisions about risk are better explained in terms of emotional affect, rather than cognitive assessment – the “risk-as-feelings” hypothesis – which seems consistent with (albeit not probative of) the notion that appetites for financial and physical risks are driven by the same emotions, and hence positively correlated.

In terms of insurance industry practices, the use of credit scores in motor and home insurance appears to be a tacit recognition of propitious selection. The practice suggests a belief in the industry that financial risk-taking and other risk-taking are positively correlated, as required by propitious selection. The Actuarial Education and Research Foundation (AERF) has recently sponsored a research project titled *Biological & psychological correlates of risk taking, credit scores and auto-losses: why insurance credit scoring works*. It is claimed that this study will show a distinct blood chemistry and financial profile for financial risk-takers.³³

²⁹ Harlow and Brown (1990).

³⁰ Barsky *et al.* (1997).

³¹ Viscusi and Hersch (2001).

³² Lowenstein *et al.* (2001).

³³ Presumably, this claim is based on Harlow and Brown (1990) as above. Details of the proposed study taken from www.actuarialfoundation.org/research_edu/r_activities.htm, accessed 28th May, 2006.

Psychological aspects of risk classification

Adverse selection stories rely on assumptions that insureds have better implicit or explicit knowledge of their risk status than the insurer, and seek to exploit this superior knowledge through insurance purchasing decisions. There are two strands of psychological research that cast doubt on these stories. The first strand is the well-known research on “decision making under uncertainty”.³⁴ This carries the implication that while people might attempt to make probabilistic decisions about buying insurance (and many other matters), they are extremely bad at doing so. The second, less well-known strand might be characterised as “decision making under ignorance”.³⁵ It carries the implication that decisions about insurance are probably not made by probabilistic thinking at all.

Decision making under uncertainty

Adverse selection stories about the interaction between individual customers and an insurer usually involve three implicit conceptual stages:

- the customer has private knowledge not available to the insurer, which is potentially relevant to their risk relative to that of other customers;
- the customer processes this information to produce better estimates of relative risk than the insurer’s; and
- the customer varies her insurance purchasing decisions in response to her relative risk estimates and the price of insurance.

These stages may be implicit rather than explicit. Individual customers need not compute superior relative risk estimates; they just need to behave in aggregate as if they do. Nevertheless, the argument that they do so behave seems predicated on at least some rudimentary or implicit form of the three stages above. But, there are psychological obstacles to each of these stages.

The first stage is the lowest hurdle. Doubtless, many customers do have private knowledge relevant to their relative risk status. For example, the life insurance customer may have a fuller knowledge of her medical history, her diet, and her physical activity level than is obtained by the insurer.

The second stage, processing this information into more accurate estimates of risk, is problematic for a number of reasons. Suppose we assume in the first instance that customer is an expert in estimating relative risk. The incremental value of the additional information may be small. Most experts are poor at making probabilistic judgments.³⁶ In experimental studies that compare expert predictions based on full knowledge with statistical models based on a limited range of factors, simple models consistently outperform experts.³⁷

³⁴ The seminal reference is Kahneman *et al.* (1982).

³⁵ For example Hogarth and Kunreuther (1995).

³⁶ Kahneman *et al.* (1982).

³⁷ In the psychological literature, this is called “clinical versus statistical prediction.” Meehl (1996) is a survey of the evidence, a revision of work first published as long ago as 1954.

More realistically, the customer is likely to have very inferior skills compared with the insurer in estimating relative risk. This reflects not only technical knowledge and abilities, but also psychological biases. For example, in motor insurance, safe driving skills appear potentially relevant to risk estimates; but results like the notorious finding that 46 per cent of U.S. drivers believed they were among the top 20 per cent safest drivers³⁸ call into question the ability of drivers to process their private information about driving skill into better relative risk estimates. In relation to life insurance, there are several studies suggesting that mortality risks are probably denied rather than over-estimated³⁹ (and this is why life insurance is sold, not bought). These biases may result in insurers making better relative risk estimates than their customers, even though the customers have additional private and relevant information.

The third stage, modifying insurance purchasing decisions in response to relative risk estimates, is also problematic. Clinicians find it notoriously difficult to persuade patients to modify their diet or lifestyle in response to relative risk estimates. Why should insurance purchasing decisions be different?

Decision making under ignorance

Theories of adverse selection implicitly assume that insurance purchasing decisions are probabilistic – that insureds think like actuaries or economists are supposed to think. Some psychological research suggests that most people probably do not think about probabilities at all, except possibly (but even then, not necessarily) when the figures are given to them.⁴⁰ Instead, they are driven largely by an affect heuristic – that is, decisions are determined by the immediate and automatic positive or negative feeling that a stimulus provokes – and then use arguments such as peace of mind or sleeping well at night to justify these decisions. If people make decisions about insurance mainly in these ways, adverse selection stories may not be a good description of real decision making.

One phrase in the last paragraph was worded carefully: “*like actuaries or economists are supposed to think*”. But I suspect that even actuaries and economists – or, more generally, the small minority of people with a probabilistic mindset – often do not think probabilistically about very unlikely events. I suggest that there is a probability threshold below which computation of expectations across different states of the world is otiose for even the most probabilistic thinker. (In deciding to accept the invitation to the conference on the theme of risk classification and public policy, how many of the audience carefully evaluated their expected utility across all states of the world – reaching Berlin, their plane crashing, etc?) This is not the same as saying that all people ignore very unlikely events all of the time (although many do, a lot of the time). Rather it is saying that decisions about insurance against very unlikely adverse events are not as responsive to price and risk changes as adverse selection stories imply.⁴¹

³⁸ Svenson (1981).

³⁹ One well-known study is Weinstein (1980) in which a large majority of people expected to live beyond 80 years.

⁴⁰ For example Hogarth and Kunreuther (1995).

⁴¹ Symmetrically, decisions about very unlikely positive events are probably also insensitive to probability

Insurers tacitly recognise that most life insurance purchasing decisions are not made probabilistically by the ways in which they approach marketing and selling of insurance. It is extremely rare for probabilistic information to be promoted to the customer in any useable manner.⁴² Instead, advertisements and sales procedures aim to engage emotions, either by emphasising the undesirability of the event insured against, or the reassurance provided by the insurance policy.

Asymmetric information – whose benefit?

Summarising the previous section, the contribution of psychological research is that it highlights how insurers could often be better informed about relative risk, even when customers possess, and attempt to use, relevant private information. Information asymmetry may sometimes be to the advantage of the insurer rather than the insured.

There is another way in which asymmetric information works to the benefit of the insurer. This is that the insurer has a vastly better understanding of the technical language of the contract, the legal context in which it will be interpreted, and the algorithms by which the premium is calculated. Insurers do not generally disclose to the customer how premiums are calculated, or the reasons why loadings have been applied. It can be very difficult for the customer to obtain this information, or indeed to establish the fact that a non-standard premium loading has been applied.

One example of how this superior knowledge is exploited was mentioned above: the use of credit scores in motor and household insurance risk classification. Customers generally do not know their credit scores, or how this might be considered relevant to motor or household insurance risks. Another example is the practice of inertia pricing,⁴³ whereby renewal premiums are varied not in accordance with risk factors, but rather in accordance with statistical predictions of the propensity of a customer with particular characteristics to “shop around” (that is, to obtain alternative quotations) at renewal time. For example, if middle-aged customers living in rural areas insuring five-year old cars have a lower propensity to “shop around” than young customers in urban areas insuring one-year old cars, the increase in premium at renewal time would be higher for the former group, after conditioning for risk factors. But when asked to explain or justify premium variations, the insurer would make reference only to risk factors. This incomplete explanation represents an information asymmetry in the insurer’s favour.

variations. The observation then offers an explanation for simultaneous participation in lotteries and purchase of insurance: both decisions are driven by affect, and hence are relatively insensitive to probability variations.

⁴² On the few occasions when probabilistic information is presented, it usually appears tendentious: it is carefully selected to evoke emotions of fear and dread (e.g., “one in three people will die from cancer”), rather than to promote accurate probabilistic thinking.

⁴³ Faculty and Institute of Actuaries (2001).

Competition in risk classification: the malevolent invisible hand?

Greater competition is a common prescription for markets that do not work well. It is generally suggested that competition between self-interested agents in markets serves the public good through the mechanism of an “invisible hand”. But the invisible hand can sometimes be malevolent: there are some activities in which competition tends to achieve nothing, or can even be deleterious. Risk classification may be one of these activities.

An unusual feature of insurance is that in addition to competing in the usual ways found in most service industries – price, level of service, recruitment of employees – insurers also compete in risk classification. That is, they compete in attracting “good risks” and avoiding (or charging higher premiums for) “bad risks”. An insurer who introduces a new form of risk classification – either charging less for some type of good risk, or charging more for some type of bad risk – can gain at least a temporary advantage over other insurers. Other insurers may then be forced to adopt the new classification, otherwise they will attract an increasing proportion of the poorer risks. So, in general, they do adopt the new classification, and the gain to the original innovator tends to be transient.

Thus competition and innovation in risk classification tends over time to lead to a progressive fragmentation of broad risk classes. When a broad risk class is fragmented as a result of insurers introducing a new risk classification, there are several effects: some customers pay less for insurance, but some pay more, and some might be excluded from insurance altogether. From a public policy perspective, it is difficult to be sure that this combination of effects is positive. It could often be negative: the reduced cost of insurance for some customers is offset by the increased cost for others, the classification process itself has costs, and the complete exclusion of some customers from insurance gives rise to problems of an “insurance underclass” of persons whom no insurer will insure. Thus competition in risk classification can lead to a worse outcome.⁴⁴

This idea may appear similar – and indeed is similar – to the conventional “adverse selection spiral” story. Adverse selection can be seen as a dual problem,⁴⁵ driven not just by the actions of insureds (as in the conventional story), but also by the competitive actions of insurers. The conventional story is usually told by those who wish to argue *against* some proposed restriction in risk classification. Their story is that if *too little* risk classification were permitted, customer decisions would lead to a spiral of increasing prices and declining coverage. Attention is focused on the deleterious effects of restrictions in risk classification. On the other hand, my argument is that *too much* risk classification – and in particular, the process of competition and innovation in risk classification – does not produce clear benefit to customers in aggregate, and can have deleterious effects. Attention is focused on the deleterious

⁴⁴ This idea appears in slightly different guise in Hoy (1984), where both higher and lower risks prefer not to see a new and only partially accurate risk classification introduced, if they do not know to which risk class they will be assigned.

⁴⁵ See Baker (2003).

effects of market competition.⁴⁶ A reconciliation of these two views is that a policymaker should target an intermediate *right amount* of risk classification. As suggested earlier, one criterion for the *right amount* of risk classification is the level which maximises loss coverage.

When does risk classification competition have public benefits?

There is one special case where competition in risk classification may be seen as unambiguously positive. This is where the risk class fragmented is the “uninsurable” class who were previously rejected by all insurers. If one innovative insurer decides to offer insurance (presumably at a high premium) to some group previously considered uninsurable, loss coverage is increased; and because the additional lives come from the uninsurable class, rather than being cherry-picked from the existing insureds, there is no effect on other insurers’ portfolios. If other insurers agree that the new classification is sound, they might follow the innovation; but this would be for a “positive” reason (the profit opportunity of additional sales) rather than the “defensive” reason, which applies when an existing insured class is fragmented. Other insurers can *choose* to follow the innovation, but there is no competitive dynamic that *compels* them to do so. This suggests that insurance of “borderline insurable” risks can be a niche business, in which continual innovation by one or several insurers need not destabilise existing risk classes.

There are some instances where risk classification competition might be seen as socially positive through signalling and incentive effects: for example, insurers’ competition to avoid insuring houses subject to flooding sends a useful signal to society about the costs of new building on flood plains; or motorists might be incentivised to take more care by accurate risk classification. But these signalling and incentive effects have limited applicability in life or health insurance. Furthermore, even in fields such as motor insurance, where incentive effects seem more likely to operate, they are in practice often difficult to discern. For example, Schwarze and Hein find that although deregulation in mandatory third party motor insurance in Germany in the 1990s led to the introduction of many new risk categories, this has *not* been associated with either a decline in total damages or a reduction in the average price of insurance.⁴⁷ In aggregate, the deregulation of risk classification in this market appears to have achieved nothing, except an increase in administrative costs.

To summarise, competition in risk classification may sometimes be guided by a malevolent invisible hand. It does not make a clear positive contribution to the aggregate welfare of customers, in the same way as competition to reduce expenses or improve levels of service. Perhaps public policy should direct competition away from risk classification, and towards more useful areas – that is, those that are more clearly beneficial to customers, for example, on operational expenses and level of service.

⁴⁶ The fact that the former possibility receives much more attention than the latter is an example of structural bias.

⁴⁷ Schwarze and Hein (2005).

An example: healthcare financing

Recent developments in healthcare financing in the U.K. can be interpreted as an application of this insight. For the past half-century, most healthcare in the U.K. has been provided by the National Health Service, which is centrally managed by government, and funded almost entirely from general taxation. Thus, there is no risk classification, and only a few small co-payments by patients for services. (Supplementary private healthcare financed by individual or group insurance does exist, but covers less than 10 per cent of the population.)

In recent years, there has been increasing recognition in government that a universal state-funded service has its own disadvantages, and that private sector disciplines including competition may have a role to play in improving healthcare. However, this increased willingness to utilise private sector concepts and disciplines has *not* involved the introduction of risk classification. Instead, private healthcare companies have been invited to bid for long-term government contracts to deliver high volumes of specific diagnostic services or clinical treatments (MRI scans, knee replacements, etc.). This approach can be interpreted as seeking to promote competition, specialisation and innovation in healthcare delivery, while avoiding competition in risk classification.

Healthcare systems in many other countries afford a much larger role to competition in risk classification, one extreme example being the United States. It is well known that the United States is a very rich country that spends a much higher proportion of GDP on healthcare than most developed countries (around 13 per cent, cf. 10 per cent in Germany, 8 per cent in the U.K.), and yet produces worse public health outcomes on a wide range of measures, including infant mortality, life expectancy at birth, and potential years of life lost.⁴⁸

For the detached observer, it is inescapable to surmise that these observations may be connected. In other words, the disappointing relative performance of the U.S. healthcare system in aggregate on many measures may be partly attributable to its relatively high reliance on competition in risk classification. There is no proof of this explanation, but it is difficult to think of a more obvious one.

Risk classification as blame

Penalising misfortune

Another perspective on risk classification is that it is a process whereby the privileged and powerful exercise judgment against the unfortunate. Risk classification can be seen as a process whereby a powerful organisation identifies and quantifies pre-existing misfortune, and then – far from seeking to help the sufferer – inflicts additional misfortune, in the form of insurance disadvantages. This is often described or justified by phrases such as “actuarial fairness” (a concept of fairness which

⁴⁸ For example see National Audit Office (2003).

non-actuarial commentators may regard as rather eccentric).⁴⁹ This allocation of additional disadvantage combined with invocations of fairness suggests the concept of risk classification as blame.

This does not attribute to proponents of risk classification the naïve view that all insurance disadvantages arise from something the insured has done, or that the insured could have avoided the disadvantage by acting differently. Rather it describes a common view about the *location and origin* of insurance disadvantage. Risk classification as blame describes the view that insurance disadvantage originates from, and is intrinsic to, the individual. For example, a disabled person's insurance disadvantage is attributed to the fact that he is disabled. This is risk classification as blame.

An alternative view is that the disabled person's disadvantage arises because society has adopted a risk classification scheme that creates and allocates the insurance disadvantage. Society could have chosen almost any scheme or level of risk classification – it could outlaw risk classification almost entirely (as in health insurance in Ireland), prohibit classification by sex (as for various products in various countries), proscribe certain genetic information, protect certain disabilities but not others, and so on. But society has chosen a scheme that allocates insurance disadvantage to this particular disabled person. In this view, the problem is not intrinsic to the disabled person; it is society's adoption of a structure which allocates disadvantage to him. The risk classification scheme – the instrument of blame – is itself the cause of disadvantage.

These contrasting views – risk classification as blame, and risk classification as the cause of disadvantage – roughly correspond to the distinction made in disability studies between the “medical model of disability” and the “social model of disability”. The medical model represents a traditional and dominant view, but most legislative and social initiatives relating to disability over the past decade or two are better understood in terms of the social model.⁵⁰

Identity, exclusion and Rawlsian veils

The discussion above was framed in respect of an unspecified “disabled person”. But suppose I had framed the discussion in terms of “a smoker”? Is the smoker's insurance disadvantage attributable to the fact that he smokes, or to society's adoption of a risk classification scheme that penalises smokers? Today, it seems few people have any hesitation in blaming smokers for smoking. But wind the clock back 30 years: the smoker would not have had an insurance disadvantage then, because risk classification made no allowance for smoking. This was not because the link between smoking and health was unknown – it had been known for 20 years.⁵¹ It was because insurers were apparently reluctant to use it in risk classification.

⁴⁹ Baker (2003) offers an alternative definition of ‘actuarial fairness’: the privileging of insurance institutions, organised on their own terms, above all other interests.

⁵⁰ Oliver (1990) provides an introduction.

⁵¹ For example Doll and Hill (1954).

This example is interesting because of the movement over a relatively short period from a situation where penalising smokers was apparently unconscionable, to the adoption of a risk classification that created a disadvantage, and thence to the apparently universal and unhesitating view that smokers are themselves to blame for this disadvantage. It highlights the variant perceptions of risk classification as a cause of disadvantage, and risk classification as blame.

Why were life insurers in the 1960s and 1970s not rushing to classify risk according to the already well-established link between smoking and mortality? It is difficult to know, but I suggest that it may have been partly for reasons of identity – to be concrete, many senior insurance executives, or their friends, their colleagues or their family, would themselves have been smokers. Risk classification as blame is not so comfortable if the blame is close to home.

This apparent reluctance to penalise smoking at a time when smoking was probably socially acceptable and commonplace in the insurance milieu leads to a more general point. One of the reasons why economists, actuaries, and other insurance insiders are often so casually comfortable with the paradigm of risk classification as blame may be that they are usually not members of the group who are blamed.⁵² It is usually some outsider – an outsider both different and emotionally distant – who is blamed, stigmatised and excluded. This disassociation of identity is consistent with the structural bias noted earlier in the paper. It is also reflected in the insiders' often tendentious terminology, such as "adverse" selection (which might not be so "adverse" to outsiders.)

Proponents of risk classification frequently assert that any penalties applied in risk classification are fair, and should be accepted as such by any observer who understands risk classification. This seems to me so obviously myopic or disingenuous that it should not require comment; but since many people seem to take it seriously, let me state the objection in the most general way possible. Imagine yourself behind a Rawlsian veil of ignorance, with a choice of being born into a world either with or without risk classification. It seems to me that any risk-averse person must choose the world without risk classification.

Another way of putting this is that if the proponents of risk classification were to become members of a stigmatised group, they might think about fairness more carefully, and they might perceive reality differently. But since they seldom are, they do not think about this.

Exclusion from debate

Distinct from the exclusions inherent in risk classification itself, another type of exclusion that often applies to broadly the same people is exclusion from debate. To give a topical example, in the U.K. there has for a number of years been debate about the policies that insurers and government should adopt relating to genetics and insurance. The terms of this debate were initially defined by insurers and actuaries, who sought to establish a mutual understanding, influence and agreement with

⁵² Of course there are exceptions, but they are rare and unrepresentative.

policymakers, medics and geneticists. In doing so, the insurers and actuaries were careful to exclude or marginalise people who might be directly disadvantaged by the genetic tests that insurers wanted to use, and the charities that represented them. These voices were never heard in debate promoted by insurers and actuaries – partly because these voices suffer various disadvantages anyway in making themselves heard, and partly because insurers and actuaries had carefully structured the debate to exclude them. Instead, they have obtained a hearing via intermittent external shocks to the insurers' and actuaries' carefully circumscribed discussion – for example, when parliamentary committees or other public bodies have sharply criticised the terms of that discussion, leading to externally imposed reforms of its structure and process.

The idea in this example has much wider application: the people who are most disadvantaged by risk classification tend to be excluded from debate about risk classification. One reason for this is the paradox of collective action. A second reason is the structural bias mentioned at the start of this paper. A third reason is that the same stigma that leads to exclusion in insurance can also devalue and discredit one's viewpoint, and can lead to practical difficulties in making that viewpoint known. If one is disabled, it may be difficult to get insurance; it may also be difficult to attend conferences about risk classification; and one's views may not be taken very seriously if one does. The same is true of poverty, extremes of age and most forms of insurance disadvantage.⁵³

Risk classification and human rights: contrary principles?

Different principles

In recent decades, risk classification practices have often been called into question by social views and laws that deprecate or prohibit discrimination according to factors such as sex, disability or race. Insurance commentators often suggest that the controversy arises because critics fail to understand the rationale of risk classification; and hence that if only insurers could better explain themselves, their practices would be accepted. However, this notion that insurers need only explain their principles for them to gain universal acceptance overlooks the fundamental and conceptual nature of the conflicts between risk classification and many contemporary social views and laws.

It is helpful to distinguish between two categories of legislation, with different conceptual bases. On the one hand, there are laws on sex and race discrimination, which are (broadly speaking) based on the concept that persons of different sex or race *should receive the same treatment*. On the other hand, there are disability discrimination laws that are based (at least in part) on the concept that persons with a disability *should receive different treatment* – in the U.K. legislation, the

⁵³ The classical sociological account of discrediting stigma is Goffman (1963). As a thought experiment, it is interesting to imagine a world in which risk classification penalises characteristics which in other contexts are associated with high social status. How much advocacy for risk classification would there be in such a world?

“*reasonable adjustment*” – focused on their disability. These two categories of legislation conflict with risk classification in different ways.

Under the first type of law, statistical discrimination is generally equally offensive as “taste-based” discrimination. The law (at least in the U.K.) draws no distinction between the employer who refuses to promote women as a matter of taste and the employer who attempts to justify his refusal by adducing statistical evidence of a correlation between sex and productivity. Statistical justification for the discriminatory practice is not regarded as sufficient grounds for permitting its use. The conflict of this type of law with risk classification is fundamental: all risk classification is statistical discrimination. It is therefore usual to make some special provisions or exceptions for risk classification in this type of law.

When insurers are faced with criticism of certain practices – such as charging different prices to men and women – they usually try to explain that this practice is based on averages, that it is not denied that many men live longer lives than the average woman, etc. In other words, insurers explain that they practise statistical discrimination. The fact that in many contexts society regards statistical discrimination *per se* as offensive helps to explain why insurers’ explanations are often not well received. The problem may not be that critics fail to understand the insurer’s principle; it may be that they do understand, and recognise it as a principle that they reject.

The conflict between risk classification and disability discrimination laws is of a different nature. In many contexts, the person with a disability may be entitled to some special treatment tailored to their disability. For example, a partially sighted customer may ask to receive a contract in large print; or a disabled traveller may ask an airline to provide a wheelchair at the airport terminal; or an employee may ask for some part of a job to be modified. Service providers and employers in the U.K. have some obligation to meet these requests (subject to a test of reasonability), and they cannot charge the disabled person for doing so. The law recognises that service providers may incur extra costs in serving a disabled person, and explicitly prohibits the charging of these costs to the disabled person. This is the whole point of the law.

In risk classification, the paradigm is entirely different: the whole point of the activity is to charge any additional costs attributable to disability to the disabled person. Proponents of risk classification often defend this as a *matter of principle*. The fact that the principle directly contradicts the principle that the law mandates for all other services helps to explain why the explanation is often not well received. Once again, the problem may not be that critics do not understand the principle; it may be that they do understand, and find it offensive.

In my view, insurers are fighting a losing battle when they seek to justify charging disabled people for the costs of disability as a *matter of principle*. Insurers may believe in the principle, but society seems to disagree, and indeed for most contexts has legislated to mandate a contradictory principle. A more robust justification for the different paradigm in risk classification as compared with other areas is that it is a *matter of degree* or a *matter of necessity*, rather than a matter of principle. Some types of disability may lead to very large additional costs in insurance; if these are large enough and relate to sufficient numbers of insureds, the whole scheme of voluntary

insurance might be endangered. Charging the costs of disability to the disabled person may be to some degree necessary to prevent this.

Adopting this justification makes it necessary to grapple with questions of degree: how much risk classification do we need to make private insurance work? Would more or less risk classification produce a better outcome? As discussed earlier in this paper, better for the insurance industry may not always be the same as better for public policy.

Malign theories

The effort that economists and actuaries direct towards propounding theories of adverse selection in insurance sometimes has an unattractive aspect. Calls to restrict the use of particular risk classifications are often motivated by altruistic human rights ideals, or concerns of equity relating to socially disadvantaged minorities. The economist or the actuary cautions that these concerns need to be balanced by efficiency concerns relating to adverse selection, and proffers an elaborate theory. The theory is enthusiastically adopted and promoted by the insurance industry as an argument against altruism or compassion. Yet the economist or the actuary is often unable to demonstrate the real existence, let alone the materiality, of the phenomena alleged in the theory. Even if he were able to do so, the economist's or actuary's role might be considered malign: enormous efforts and ingenuity are directed towards highlighting minor differences in efficiency as arguments against altruism or compassion. Perhaps society would prefer not to know.

Adverse selection and large policies: the one-shot gambler

It is often suggested that if insurers are banned from asking about certain underwriting information – for example, about the results of predictive genetic tests – a customer with private knowledge of their risk status will rationally exploit this information by buying a very large life insurance policy, or multiple policies. I think this idea is misguided, because the customer cannot make the “favourable” bet 1,000 times: he is a “one-shot gambler” who can make the bet only once. When one considers plausible probabilities and premia, it seems to me that despite “favourable” odds, a large bet on over-insurance is still unattractive to a person who can make the bet only once.

For example, suppose that an insurer offers me a life insurance premium based on an assumed risk of $p=1$ per cent, but I have private information that tells me my real risk is $p^*=4$ per cent. This information might make me worry more about my dependants, and make me more likely to buy life insurance. But should I, as many commentators seem to suggest, go further and buy a policy (or a number of policies) with unusually large sum assured, say €1 million or €10 million?

If I could make a bet of this nature on 1,000 independent lives concurrently, then it might be attractive. But, in reality, I can only make it on one life. I would also need to pay the very large premium, and it is almost certain (96 per cent certain, on the true probabilities) that I just lose the premium. This seems to me a very unattractive proposition. It would remain unattractive for a wide range of plausible values of the

premium probability p and the true probability p^* . In my view, the “one-shot gamble” starts to become attractive only if the insured event approaches a certainty ($p^* > 75$ per cent, say).⁵⁴

The unattractiveness and implausibility of “one-shot gambles” in most realistic insurance scenarios is reflected in the paradigm adopted by insurance sales agents and advisers. The paradigm is about identifying the customer’s insurance needs, and then recommending policies that meet those needs in a cost-effective way. The customer’s needs (family responsibilities, mortgage debts, etc.) dominate the decision. The paradigm is not concerned with maximising the customer’s wealth by exploiting small variations in insurance prices. Price/risk mis-pricings for particular risks can to some degree persist, because they are not exploitable by a one-shot gambler. Mis-pricings in insurance are potentially exploitable *by other insurers* through risk selection competition, but they are not easily exploitable *by customers*.

Conclusion

This paper has presented a number of novel – or at least, seldom-aired – perspectives on risk classification. Many of its observations are rather obvious, and some may not be controversial; but they tend to be seldom made. The key ideas are as follows:

- Adverse selection may not always be adverse. From a public policy viewpoint, some degree of adverse selection offers the advantage of “insurance bought by people likely to need it”.
- A public policymaker concerned with risk classification should consider the criterion of loss coverage. If insurance is viewed as a social good, the socially optimal level of adverse selection is probably the level which maximises loss coverage. Additional or alternative policy objectives are possible.
- The commercial objectives of an insurer may be different from those of the public policymaker, and might be well served by *minimising* loss coverage: “insurance sold to people unlikely to need it.”
- Theoretical models of risk classification may exaggerate the significance of adverse selection. In econometric studies, adverse selection is often difficult to discern. As a concept, adverse selection may be more apposite in explaining interactions between competing insurers offering close substitutes, rather than interactions between insurer and insureds.
- Propitious selection arises where insurance purchase is associated with lower risk, not higher risk. There are economic and psychological literatures supporting this idea, and some insurance industry practices that tacitly acknowledge it.
- Psychological perspectives on decision making under uncertainty, or decision making under ignorance, can help explain why the phenomena predicted by adverse selection models are often difficult to observe.

⁵⁴ Another insight into this is to consider the putative bet on over-insurance as part of the policyholder’s overall portfolio of assets. The bet on insurance has a small positive edge, but a very high chance of loss; the optimal portfolio allocation is therefore very low.

- Risk classification competition can be guided by a malevolent invisible hand: it probably produces less customer benefit in aggregate than other forms of competition. A policymaker should try to focus competition in more useful areas.
- Risk classification often appears to involve penalising the unfortunate for their misfortune. The common assertion that this is fair seems myopic or disingenuous: behind a Rawlsian veil, any risk-averse person would prefer to be born into a society with no risk classification.
- Risk classification is based on principles of (a) statistical discrimination and (b) charging the costs of disability to a disabled person. These principles are often in contradiction to (a) sex and race discrimination laws and (b) disability discrimination laws, respectively.
- The fallacy of the one-shot gambler: suggestions that people with private knowledge of high-risk status will rationally exploit insurers by effecting very large policies confuses the position of insurance companies, which make many uncorrelated bets, with that of the customer who can bet on only one event. For a customer who can bet only once, the idea of “taking advantage” of the insurer by effecting large amounts of insurance against a single very unlikely event is arithmetically implausible.

As noted in the introduction, much discussion of risk classification and public policy tends to be affected by a structural bias, in that it is funded or influenced by the insurance industry, or by researchers whose advancement is best served by conforming to an industry perspective. In academic discussion of insurance there is often another structural bias. This is the bias arising from the predominance of economists, who proceed by assuming that people are invariably selfish maximisers – that the world is full of rational fools, in Amartya Sen’s famous phrase.⁵⁵ Notoriously, a good deal of work in experimental economics has suggested that this assumption might be more valid for economists themselves than for non-economists.⁵⁶ Perhaps this partly explains the gulf between what economists and actuaries typically think about fair and unfair discrimination in insurance, and what many other people think.

Acknowledgements

I thank Philip Cooper, workshop participants at DIW Berlin in June 2005, and two anonymous referees for comments on this paper.

References

- Baker, T. (2003) ‘Containing the promise of insurance: adverse selection and risk classification’, *Connecticut Insurance Law Journal* 9: 371–396.
- Barsky, R.B., Juster, F.T., Kimball, M.S. and Shapiro, M.D. (1997) ‘Preferences parameters and behavioral heterogeneity: an experimental approach in the health and retirement study’, *Quarterly Journal of Economics* 112: 537–579.

⁵⁵ Sen (1977).

⁵⁶ For a survey, see Frank *et al.* (1993).

- Blows, S., Ivers, R.Q., Connor, J., Ameratunga, S. and Norton, R. (2003) 'Car insurance and the risk of car crash injury', *Accident Analysis and Prevention* **35**: 987–990.
- Cardon, J.H. and Hendel, I. (2001) 'Asymmetric information in health insurance: evidence from the National Medical Expenditure Survey', *Rand Journal of Economics* **32**: 408–427.
- Cawley, J. and Philipson, T. (1999) 'An empirical examination of information barriers to trade in insurance', *American Economic Review* **89**: 827–846.
- Chiappori, P.-A. and Salanie, B. (2000) 'Testing for asymmetric information in insurance markets', *Journal of Political Economy* **108**: 56–78.
- Crocker, K.J. and Snow, A. (1986) 'The efficiency effects of categorical discrimination in the insurance industry', *The Journal of Political Economy* **94**: 321–344.
- Doll, R. and Hill, A.B. (1954) 'The mortality of doctors in relation to their smoking habits', *British Medical Journal* **228**: 1451–1455.
- Dyson, R. (2006) '1 in 5 critical illness claims thrown out', *Mail on Sunday*, 27th March, 2006.
- Eassie, Lord (2006) Judgment in the case of *Cuthbertson v Friends Provident Life Office*, Court of Session, May 2006, from www.scotcourts.gov.uk/opinions/2006CSOH74.html, accessed 1st August, 2006.
- Faculty and Institute of Actuaries (2001) *Inertia pricing – actuarial profession briefing statement*, from http://www.actuaries.org.uk/Display_Page.cgi?url=/news/inertia_pos.xml, accessed 1st August, 2006.
- Finkelstein, A. and Poterba, J.M. (2000) *Adverse selection in insurance markets: policyholder evidence from the UK annuity market*, NBER Working Paper, 8045.
- Frank, R.H., Gilovich, T. and Regan, D.T. (1993) 'Does studying economics inhibit cooperation?', *Journal of Economic Perspectives* **7**: 159–171.
- Goffman, E. (1963) *Stigma: Notes on the Management of Spoiled Identity*, New York: Simon and Schuster.
- Harlow, W.V. and Brown, K.C. (1990) 'Understanding and assessing financial risk tolerance: a biological perspective', *Financial Analysts Journal* **46**: 50–62.
- Hemenway, D. (1990) 'Propitious selection', *Quarterly Journal of Economics* **105**: 1063–1069.
- Hemenway, D. (1992) 'Propitious selection', *Journal of Risk and Uncertainty* **5**: 247–251.
- Hogarth, R. and Kunreuther, H. (1995) 'Decision making under ignorance: arguing with yourself', *Journal of Risk and Uncertainty* **10**: 1015–1036.
- Hoy, M. (1984) 'The impact of imperfectly categorising risks on income inequality and social welfare', *Canadian Journal of Economics* **17**: 557–568.
- Human Genetics Commission (2001) *Whose hands on your genes?*, London: Department of Health.
- Hunt, J. and Birnbaum, B. (2001) *Credit Insurance Overcharges Hit \$2.5bn Annually*, Washington, DC: Consumer Federation of America.
- Inman, P. (2004) 'Barclays insurance hideaway', *The Guardian*, 8th March, 2004.
- Kahneman, D., Slovic, P. and Tversky, A. (1982) *Judgments Under Uncertainty: Heuristics and Biases*, New York: Cambridge University Press.
- Lowenstein, G.F., Hsee, C.K., Weber, E.U. and Welch, N. (2001) 'Risk as feelings', *Psychological Bulletin* **127**: 267–286.
- Meehl, P.E. (1996) *Clinical versus Statistical Prediction: A Theoretical Analysis and a Review of the Evidence*, Northvale, NJ: Jason Aronson.
- National Audit Office (2003) *International Healthcare Comparisons*, London: National Audit Office.
- Oliver, M. (1990) *The Politics of Disablement*, London: Macmillan.
- Rothschild, M. and Stiglitz, J. (1976) 'Equilibrium in competitive insurance markets: an essay on the economics of imperfect information', *Quarterly Journal of Economics* **90**: 629–649.
- Schwarze, R. and Hein, T. (2005) *Is the market classification of risk always efficient? Evidence from German third party motor insurance*, CARR Working Paper, 32.
- Sen, A. (1977) 'Rational fools', *Philosophy and Public Affairs* **6**: 317–344.
- Siegelman, P. (2004) 'Adverse selection in insurance markets: an exaggerated threat', *Yale Law Journal* **113**: 1225–1271.
- Svenson, O. (1981) 'Are we all less risky and more skilful than our fellow drivers?', *Acta Psychologica* **47**: 143–148.

- Thomas, R.G. (2001) 'Genetics and insurance – an actuarial perspective with a difference', *Response to consultation by Human Genetics Commission London*, from <http://www.guythomas.org.uk>, accessed 1st August, 2006.
- Viscusi, W.K. and Hersch, J. (2001) 'Cigarette smokers as job risk takers', *Review of Economics and Statistics* **83**: 269–280.
- Weinstein, N.D. (1980) 'Unrealistic optimism about future life events', *Journal of Personality and Social Psychology* **39**: 806–820.
- Wilson, C. (1977) 'A model of insurance with incomplete information', *Journal of Economic Theory* **16**: 167–207.

About the Author

R. Guy Thomas is an actuary and investor, and formerly Head of Actuarial Science at the University of Kent.