

Reinsurance for Catastrophes and Cataclysms

David M. Cutler and Richard J. Zeckhauser

Insurance has done more than all gifts of impulsive charity to foster a sense of human brotherhood and of common interests. It has done more than all repressive legislation to destroy the gambling spirit. It is impossible to conceive of our civilization in its full vigour and progressive power without this principle, which unites the fundamental law of practical economy, that he best serves humanity who best serves himself, with the golden rule of religion, "Bear ye one another's burdens."

—William Gow, "Insurance," *Encyclopedia Britannica* (11th ed., 1910)

This paper examines the optimal design of reinsurance—insurance for insurers. Reinsurance enables insurers to lay off concentrated positions in idiosyncratic risks and places risks that are large for society as a whole in the hands of those best able to bear them. We examine how reinsurance can be designed to meet these purposes. Although our focus is on reinsurance, much of what we discuss is widely relevant to primary-insurance markets as well.

Reinsurance is valuable for insurers because it allows them to reduce their risk levels. If an insurance company is concentrated in a risk in a particular geographic area, for example, then a disaster affecting everyone in that area may severely reduce the insurance company's reserves. We term a situation of threatened or actual insurer insolvency in the event of an adverse risk a *catastrophe*.

The first question that we address is, How should reinsurance premiums and reimbursement be designed to address the potential for catastrophes? This question is complicated because the primary insurer is likely to know much more about the true loss distribution than will the reinsurer. As a result, adverse selection is a concern. We show that optimal policies in the face of adverse selection depend on the nature of the information asymmetry. When the asym-

David M. Cutler is professor of economics at Harvard University and a faculty research fellow of the National Bureau of Economic Research. Richard J. Zeckhauser is the Frank P. Ramsey Professor of Political Economy at the Kennedy School of Government, Harvard University, and a research associate of the National Bureau of Economic Research.

The authors are grateful to Roy Astrachan and Sandra Favlukas for research assistance; to David Cummins, Christopher Lewis, and Richard Phillips for making available their value-adjusted loss data; and to Chris Avery, John Cochrane, Ed Glaeser, Bill Hogan, Louis Kaplow, Jeff Liebman, John Pratt, Aaron Stern, Ray Vernon, and both academic and industry conference participants for helpful discussions. The authors acknowledge partial financial support for this paper from the National Institutes on Aging, the Herrnstein Fund, and the NBER Project on the Financing of Property/Casualty Risks.

metry is on the probability of a loss but not its magnitude, the optimal reinsurance contract is a standard “excess-of-loss” policy—the primary insurer is responsible for small risks, and the reinsurer is responsible for large risks. When the asymmetry is over the magnitude of the loss, the optimal reinsurance policy covers smaller losses as well as large losses but leaves the primary insurer exposed for some large losses. The striking difference between these policies suggests that one cannot make broad statements about the optimal form of reinsurance policies.

As the size of catastrophes increases, a second difficulty with reinsurance markets emerges. If the underlying risk is large in aggregate, the reinsurance industry as a whole may not be able to provide sufficient capital to cover a loss. Hence, wide-scale insolvencies may be threatened, and insurance may become unavailable or excessively priced. We term an event that strains worldwide insurance and reinsurance industry reserves—arbitrarily defined as an event with \$5 billion or more of insured losses—a *cataclysm*.¹

The second question that we address is, Who should provide reinsurance for cataclysmic events? By definition, the reinsurance industry will be poor at reinsuring cataclysms. We argue that cataclysms must be reinsured either in broader financial markets or by the government, or both. The second part of the paper explores these options.

The paper is structured as follows. We begin in the next section by outlining the nature of the problem. In the second section, we consider quantitatively the size of property-casualty insurance and reinsurance markets. In the third section, we deal with the problem of reinsuring catastrophes. In the fourth section, we consider reinsurance of cataclysms. The last section concludes.

6.1 The Nature of the Problem

To understand the nature of our analysis, consider a particular example: the market for homeowner’s insurance. Fire damage is one risk facing insurers underwriting homeowner’s policies. If the insurer is sufficiently big, given that the risks being insured are for the most part uncorrelated, internal diversification will provide sufficient diversification for this risk. Reinsurance is internal.

Other risks will not be diversifiable in this fashion, however. For example, a large hurricane or earthquake will affect all the houses in a broad geographic area. If insurers are concentrated in particular areas, for example, because of differences in sales-force concentration or underwriting knowledge,² they may be overly exposed to particular risks. Thus, insurers concentrating in California may find themselves with excessive amounts of earthquake risk, while insurers predominantly in Florida may be overly exposed to hurricane risk. Adding ad-

1. More generally, a cataclysm occurs when the size of a loss is large relative to the insurance pool for that category of risk.

2. With insurance, as opposed to most other goods, characteristics of the purchaser matter. Hence, knowledge of local conditions provides for local economies of scale.

ditional houses in California or Florida to the pool will not help diversify risk if an earthquake or a hurricane will damage all houses.

This excessive concentration will be undesirable for the insurance company's shareholders if the firm is a stock company and the shareholders' portfolios are not otherwise diversified,³ or for its policyholders if the firm is a mutual. Such concentration may even risk the insurer's bankruptcy, putting recovery for its insureds at risk.⁴

Reinsurance can help insurers diversify their excessive concentration in particular markets. For example, insurers in California and Florida can swap earthquake and hurricane risk, or both sets of insurers could sell risks to third companies, including specialty reinsurers, perhaps in other countries. We term this role for reinsurance the *catastrophe role*. A catastrophe is bigger than any one insurer can handle but not big enough to upset the entire insurance market. A catastrophe is defined as an event with at least \$5 million in insured damages that affects the policyholders of many different insurance companies (Insurance Information Institute 1996a).⁵ A \$10 million loss at one company, for example, may be large for that company, but for the industry as a whole it is a minor event. Pooling of catastrophic losses is the traditional role for reinsurance.

A major difficulty with catastrophic reinsurance is adverse selection. An insurer who is buying risk protection for his portfolio will have superior information about the risks he holds, partly because he was more knowledgeable at the time the risk was written and partly because of his subsequent experience. If diversification is the goal, the ideal seller would be a company whose own book is quite different in its holdings;⁶ naturally, such a seller is likely to be poorly informed about true risk levels. Both seller ignorance and buyer superiority of information suggest that adverse selection will be a problem (Akerlof 1970).

Moral hazard may also be a concern with reinsurance. Insurers that have purchased reinsurance may have differing incentives to sign up particular insureds, monitor the precautions that their insureds take, or fight potential litigation. We leave these issues of moral hazard for later analysis.

Until recent years, adverse selection was compounded because reinsurance was transacted more on a relationship basis—insurers and reinsurers dealt frequently on a business and social level—than on a scientific basis. In such a market, it is easy to see how private information may be exploited. Reinsurers

3. If shareholders have well-diversified portfolios, a good deal of risk spreading can be achieved by having companies bear the risk of losses and passing them on, in turn, to their shareholders. We suspect that this is how a large amount of risk is ultimately diversified.

4. Insurance managers are in an agency relationship with policyholders. Bankruptcy in response to a massive loss is probably a larger relative loss for policyholders. If, as seems likely, policyholders cannot monitor sufficiently, insufficient reinsurance purchases may result. For related reasons, states impose solvency requirements on insurers.

5. The loss threshold for a catastrophe was raised to \$25 million in 1997.

6. This is why reinsurance at the international level is an important phenomenon.

today rely much more on statistical methods, risk models, etc.⁷ No matter how hard reinsurers try to learn about the relevant information, however, adverse selection remains a concern. Consider an extreme case—where reinsurers know more about the underlying risk being insured than does the primary insurer. Suppose that the risk being insured is commercial liability. The insurer may still have relevant private information—how likely the particular company being underwritten is to file for a claim; how aggressive it will be in legal actions; how competent the executives are; etc. Reinsurer information will be reflected in competitive market prices for the risks being insured, but, within that risk class, private information possessed by the primary insurer may still influence when reinsurance is purchased.

Often, the risk being insured is so large that the aggregate risk exposure of the worldwide insurance industry is dangerously high, even if risks were fully reinsured. Hurricane Andrew, the largest U.S. catastrophe on record, cost insurers \$15.5 billion. Had Andrew hit Miami rather than less densely populated parts of South Florida, the losses could have been at least three times as large. The second largest catastrophe on record was the Northridge Earthquake of 1994, with insured losses of about \$12.5 billion. Had the earthquake hit Los Angeles or San Francisco, the damage could well have been over \$50 billion. If a major earthquake hit Tokyo, the losses could approach \$1 trillion. Beyond natural disasters, the cost to insurers of toxic waste cleanup may reach as high as \$30–\$50 billion (Insurance Information Institute 1996b). The potential for cataclysms to strike insurance markets is very real.

When cataclysms strike, insurers may be undercapitalized, and widespread insolvencies may result. To avoid such a situation, insurers might decide not to underwrite particular risks. For example, some insurers—such as Nationwide—no longer write hurricane insurance in Florida. Others may refuse to cover various environmental, product safety, or other risks that stretch years into the future (long-tailed risks). When terrorist bombings struck London in 1992, for example, reinsurers and insurers immediately stopped insuring terrorism damage. Less drastic measures to curtail exposure to cataclysms would be to switch from occurrence to claims-made policies (from year of the accident to year of the lawsuit) and to write policies for short periods of time. The net result of such actions would be to limit or dry up reinsurance markets, which in turn would dry up insurance and thus stifle economic activity.

Why do insurers and reinsurers not merely raise rates when new risks emerge or old risks appear to be greater in magnitude? There are two answers. First, implicit and explicit regulatory constraints may prevent this. Second, insurers appear to be extremely risk averse about open-ended risks, and the risk premium required for them to continue bearing this risk would therefore be prohibitively expensive.

7. We are grateful to participants at the conference for discussing the actual operation of reinsurance markets with us.

By definition, the reinsurance market cannot provide coverage for cataclysms. Reinsuring such risks requires new institutional forms. While \$10 billion is large relative to the insurance markets, it is tiny relative to asset markets as a whole. The total value of stocks traded on the New York Stock Exchange, for example, is over \$6 trillion, and the aggregate value of assets in capital markets in the United States is estimated at more than \$10 trillion. Thus, one possibility would be to diversify cataclysms through broader capital markets, and some recent steps have been taken in that direction. Cataclysms can also be dealt with through government reinsurance—and have been in many nations. The government has the advantage that it can diversify losses over time and over broader groups of people than the investors in asset markets. Section 6.4 below explores the role of reinsurance for cataclysms, looking in particular at publicly provided reinsurance.

6.2 Reinsurance Markets and Potential Losses

Before considering the design of reinsurance policies, we start with a discussion of reinsurance markets and the potential for cataclysmic losses.

6.2.1 The Size of Reinsurance Markets

Property-casualty insurance—ranging from fire insurance to products liability—is a \$300 billion per year industry. Reinsurance is a much smaller but still significant feature of property-casualty insurance. In 1994, about \$140 billion of insurance was reinsured to an affiliate of the primary insurer, and \$47 billion was reinsured to nonaffiliates (A. M. Best 1996).

Reinsurance is written by reinsurance departments of primary insurers as well as specialty reinsurers (the most famous of which is Lloyd's of London). Specialty reinsurers account for about \$12 billion of premium income annually. While primary insurers tend to write primarily in their own country, reinsurance is a much more international business. Nearly half the reinsurance ceded by U.S. insurers to nonaffiliates is ceded to international reinsurers. The leading reinsurance markets are the United States, London, and, increasingly, Bermuda (Giles 1994; Insurance Information Institute 1996c).

Table 6.1 shows which lines of insurance are reinsured most heavily. The first column shows direct premiums written. The second column shows all premium cedings to both affiliates and nonaffiliates.⁸ The third column shows cedings to nonaffiliates only. The largest property-casualty categories are auto liability (direct premiums of \$81 billion) and automobile physical damage (direct premiums of \$43 billion). Reinsurance for automobile losses is comparatively rare; less than 10 percent of the premiums are ceded to nonaffiliates.

8. The amount of reinsurance can be greater than direct insurance written since risk may be ceded to affiliates and then further ceded to nonaffiliates. This transaction would be counted twice in the table.

Table 6.1 Size of the Property-Casualty Insurance and Reinsurance Industries, 1995

Line of Business	Direct Business (\$ billion)	Reinsurance Ceded as a Percentage of Direct Business	
		Affiliates and Nonaffiliates	Nonaffiliates Only
Automobile liability	81	54	9
Automobile physical damage	43	48	8
Workers' compensation	28	88	16
Home/farm multiple peril	27	61	11
Products and other liability	25	89	30
Commercial multiple peril	21	86	16
Special property ^a	19	98	45
Other ^b	11	60	18
Medical malpractice	6	42	22
Special liability ^c	4	112	53
Fidelity/surety	4	64	21
Financial/mortgage guaranty	2	22	11
Total	273	67	16

Source: A. M. Best.

^aFire, allied lines, inland marine, earthquake, glass, burglary, and theft.

^bCredit, accident and health, international, other.

^cOcean marine, aircraft (all perils), boiler and machinery.

This makes sense; automobile losses result from thousands of independent accidents, so diversification is accomplished internally. The one major risk associated with automobile accidents is that the litigation climate will change. Under our present system, where premiums are set annually, this risk is primarily borne by drivers, who incur any yearly variability in expected losses. Reinsurance is most common for special liability (ocean marine, aircraft, boiler and machinery), special property (fire, earthquake, etc.), and products and other liability. With some exceptions, these risks tend to be more correlated across individuals or over time than is automobile liability.

Most reinsurance is written on a "treaty" basis—reinsurance is provided for all exposures for a specific class or multiple classes of business.⁹ Treaty reinsurance is generally of two types. The first type is *proportional reinsurance*; the reinsurer agrees to assume a share of the risk in exchange for a share of the premiums. Typically, the ceding company retains 30 percent or less of the original risk. The second type of reinsurance is *excess-of-loss reinsurance*; the ceding company retains risk up to some amount, and then reinsurance pays above that. Excess-of-loss reinsurance is often purchased in layers. For ex-

9. For more discussion of reinsurance policies, see Conning & Co. (1993).

ample, a company might purchase a layer of reinsurance for \$1.5 million in excess of \$500,000—that is, the reinsurer assumes risks between \$500,000 and \$2 million. Excess-of-loss reinsurance has a greater orientation toward catastrophe protection than does proportional reinsurance. We analyze both these forms of insurance below.

6.2.2 The Potential for Cataclysms

We term events that would exhaust the reserves of worldwide insurers in a particular risk market *cataclysms*. We arbitrarily call a loss *cataclysmic* if it exceeds \$5 billion. Much of the concern about reinsurance markets is related to the potential for cataclysmic losses. It is useful to understand how these potential losses might arise and what their magnitudes might be.

Cataclysmic losses can arise in two ways. The first is “single-event” cataclysms, such as a major earthquake in San Francisco or Tokyo that imposes sufficiently large losses to exhaust insurer capital. The second is what we label a *common-risk* cataclysm—a relatively small adverse change in a risk that affects large numbers of people and thus imposes large expected losses on insurers as a whole. Changes in climatic conditions that increase the probability of a hurricane but not the damage per hurricane are an example of a common-risk cataclysm; increases in tort-liability judgments for a given accident size and type are another.¹⁰ With either type of cataclysm, unfolding events may provide significant information about underlying, possibly changing, risk levels. If so, as we shall see in section 6.4 below, risks about future premiums levels may be significant.

Single-Event Cataclysms

We consider the potential for single-event cataclysms using data from 1949 to 1994 compiled by Property Claim Services (PCS). Damages should naturally increase over time as the value of the asset being insured—generally house prices—rise. To express losses on a consistent basis over time, we follow Cummins, Lewis, and Phillips (chap. 3 in this volume) and inflate losses to 1994 dollars using the growth of aggregate house values in each state and year.¹¹

PCS has raised the threshold for a catastrophe over time; we cannot adjust for this, but we suspect that threshold changes do not have a major effect on the time series of losses. Cataclysmic losses are dominated by several large events. In 1992, losses were almost \$25 billion, largely because of Hurricane Andrew. Losses were nearly \$20 billion in 1994, with a substantial share from the Northridge Earthquake. Losses were nearly \$10 billion in 1989 (Hurricane Hugo).

10. Common risk from changing liability rules is sometimes referred to as *social inflation*.

11. Updating by aggregate house values factors in three terms: population growth, overall price inflation, and the growth of real house values. We are grateful to David Cummins, Christopher Lewis, and Richard Phillips for supplying us with their data.

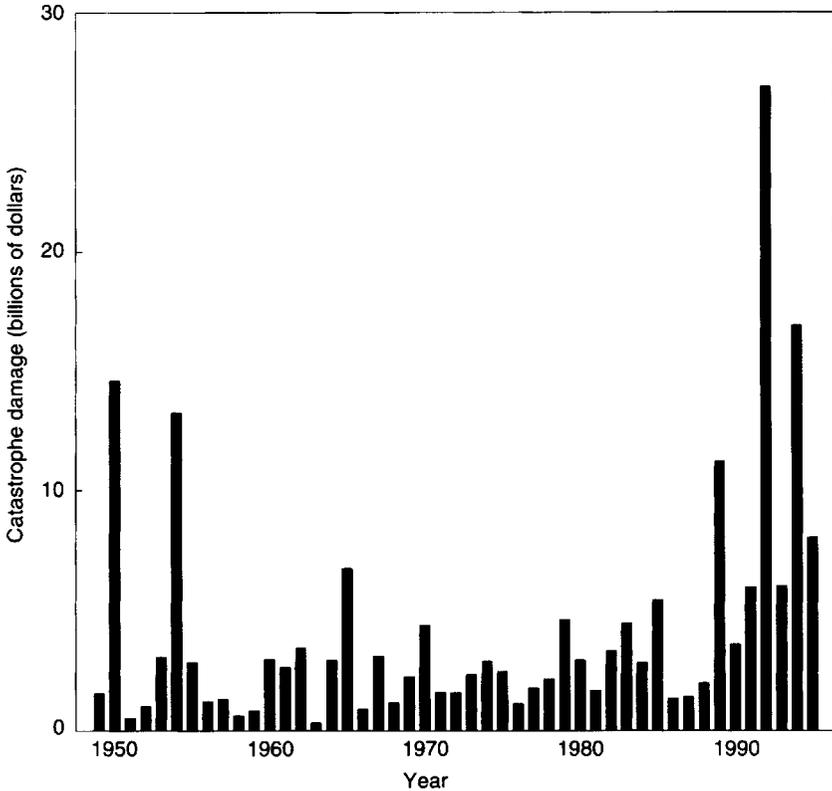


Fig. 6.1 Catastrophic/cataclysmic losses over time

Note: Data are from Property Claim Services as adjusted by Cummins, Lewis, and Phillips (chap. 3 in this volume).

Cataclysmic damages have increased over time. There were two large losses in the 1950s (windstorm damage in the Northeast and Middle Atlantic in 1950 and two hurricanes hitting the Northeast and Middle Atlantic in 1954) but no losses of even \$7 billion between 1955 and 1988. Since 1989, however, cataclysmic losses have been much larger. On the basis of figure 6.1, an insurer thinking about pooling cataclysmic losses over many years could not be confident of setting appropriate premiums since past losses do not seem to be a reliable indicator of future losses.

The nature of the risk greatly affects the potential for cataclysms. We compare losses from two types of events—windstorms and hurricanes.¹² Wind-

12. Data on losses by catastrophe are available only through 1994. In some years, there are no catastrophic losses from hurricanes. We assume that catastrophe losses for hurricanes were \$500,000 in those years.

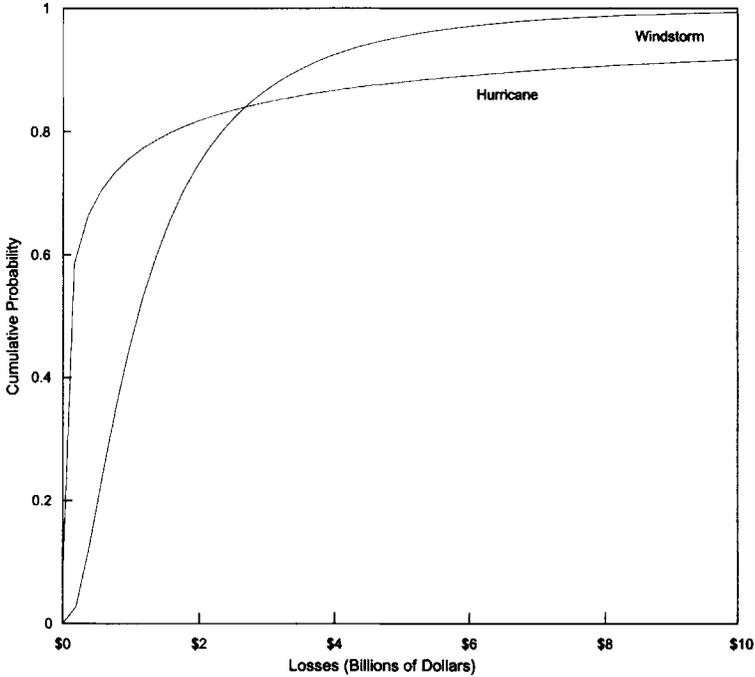


Fig. 6.2 Cumulative probability of damage from windstorms and hurricanes

Note: Data are from Property Claim Services as adjusted by Cummins, Lewis, and Phillips (chap. 3 in this volume).

storm damage averages \$1.2 billion annually; hurricane damage is substantially less, with an average of \$97 million annually. The variance of hurricane damage is much greater, however. The standard deviation of windstorm damage is 90 percent, compared to 340 percent for hurricane damage.

To determine what these different variances imply, figure 6.2 graphs the cumulative probability of loss for each type of risk on the basis of a lognormal model of expected losses.¹³ For windstorms, the probability of only minimal damage (less than \$500 million in aggregate) is only 25 percent, and the probability of a \$10 billion loss—twice our cataclysmic level—is about 0.5 percent. In contrast, the probability of minimal hurricane damage is nearly two-thirds, while the probability of \$10 billion of losses is about 8 percent. Indeed, the loss distribution from hurricanes has an extremely wide tail; the probability of a \$20 billion hurricane, for example, is over 5 percent. Thus, while cataclysmic

13. Clearly, we have very little information about major cataclysms, so the lognormal might not be the best approximation to the true distribution in the tails, which is the primary concern. Cummins, Lewis, and Phillips (chap. 3 in this volume) show that the lognormal model fits the tails of the distribution nearly as well as do more flexible distributions.

windstorm damage does not appear to be a major concern, the possibility of a cataclysmic hurricane is very real.

Common-Risk Cataclysms

The probability of common-risk cataclysms is more difficult to judge, in part because common-risk cataclysms may result from a number of small events rather than one large event. For example, knowledge about which chemicals are dangerous is formed only over time. Similarly, legal rulings regarding insurer liability for hazardous products do not change all at once but are refined gradually in a series of judicial and appellate decisions. The cumulative effect of such small events can be very large. Indeed, it has been widely argued that legal changes in the 1970s and 1980s represented a cataclysmic change (see, e.g., Huber 1988; Viscusi et al. 1993).¹⁴ The likelihood of a common-risk cataclysm is also hard to project because common-risk cataclysms tend to be unique events. We do not expect new liability revolutions such as we have seen in recent decades.

Not all common-risk cataclysms develop slowly, however. For example, when the first terrorist incident happened in London in 1992, insurers and reinsurers were quick to perceive this as a major area of potential loss.

Often, common-risk cataclysms evolve over time, with a significant component of risk relating to the insurance coverage itself. For example, an insurance company might insure any malpractice claim in the future filed against its insured physicians in the policy year—termed an *occurrence-based* policy. Or insurers could underwrite protection from the financial consequences of leakage over the next twenty years from hazardous waste disposed this year. These types of long-tailed risks are particularly susceptible to the arrival of damaging new information or changes in legal and judicial interpretations. For example, we may learn that products previously thought safe actually damage humans—as with asbestos or the Dalkon shield—or are effectively treated that way by the legal system.¹⁵ Or legal rulings could expand an insurer's liability beyond what it thought it had committed to—as with environmental liability.

It is not even clear that these long-term risks can be measured by insurers, let alone priced. *Risk*, as it is conventionally defined, assumes that insurers know the probability of a loss and its expected size. As the period of insurance

14. It is important for policy purposes to differentiate between two factors. One thing that we learned over recent decades is that many things once thought safe—e.g., asbestos in buildings—are not in fact safe. Coupled with that is the assignment of liability once we have discovered damages. For insurance losses, the distinction between these two is irrelevant. The distinction is important for public policy purposes, however, because, while assigning risk *ex post* is just a transfer from one person to another, discovering a new set of toxic chemicals is a net cost to society that *someone* must pay. By definition, one group of people or another will be worse off.

15. For a discussion of legal and scientific evaluations of the safety of silicone-gel breast implants, see Angell (1996). Despite the lack of scientific evidence of risk from breast implants, there are likely to be multi-billion-dollar recoveries.

coverage increases, however, each of these values becomes more uncertain. Zeckhauser (1991) distinguishes between *risk*, the common basis for insurance, Knightian *uncertainty*, in which insurers know the events they are insuring but not their probability, and *ignorance*, in which even the events themselves are not well defined.

Fire insurance is overwhelmingly risk based. An example of Knightian uncertainty might be insurance against workplace accidents. The fact that workplace accidents may occur is well known, but the potential medical cost for workplace accidents several years in the future is not well known. Thus, the costs associated with work accidents several years in the future are uncertain. The potential for legal or legislative changes in liability is an example of ignorance. The provisions of the Superfund legislation, for example, are not something that insurers in the 1950s and 1960s could reasonably have had much idea about. Insurers are also ignorant about potential losses from terrorist risk. (In the life-insurance domain, we are ignorant about the risks of a future AIDS-type epidemic.)

Insurers and reinsurers are reluctant to provide protection against uncertainties and extremely reluctant to do so against events about which they are ignorant. To provide coverage in such a situation would entail high risk and would allow for significant adverse selection since others could know much more. Beyond this, the decision maker—say a manager as agent for shareholders—would be open to blame from shareholders: “Knowing nothing about it, how could you write insurance against that type of event?”

Changes in common risk can be significant when the time between when the losses occur and when the damages are paid is greater. This is a period when additional knowledge about true risks is learned and expected damage awards can change the most. The time lag varies substantially for different lines of insurance. Table 6.2 shows the share of losses incurred in 1986 that were paid within a given amount of time for different lines of property-casualty insurance. (Since data do not exist past 1995, we assume that all claims are paid within nine years.) For example, the first row shows that, for home/farm and commercial multiple-peril policies, 75 percent of the claims that are paid within the first nine years of the policy are paid within the first year (by the end of 1987) and 96 percent of the claims are paid within five years (the end of 1991). These claims tend to be paid very quickly; common-risk problems may thus not be particularly severe.

For other types of risks, however, there is a much longer time between loss and recovery, creating greater risk. For example, only half of workers' compensation claims are paid within the first year and only 91 percent within five years. Falling in between home/farm and commercial insurance and workers' compensation are special liabilities (aircraft, ocean marine, and the like) and automobile liability. As a rough rule of thumb, losses involving personal injury impose a much greater common-risk problem than losses involving damage to

Table 6.2 Cumulative Share of Losses by Payment Period (%)

Line of Business	Years		
	1	5	9
Home/farm/commercial multiple peril	75	96	100
Special liabilities	63	95	100
Automobile liability	61	97	100
Workers' compensation	50	91	100
Medical malpractice, occurrence	6	67	100
Medical malpractice, claims made	20	85	100
Products liability, occurrence	9	74	100
Products liability, claims made	20	94	100
Other liability, occurrence	21	81	100
Other liability, claims made	18	85	100

Source: A. M. Best.

Note: Data are for incurred losses in 1986. Year 1 extends through 1987; year 5 extends through 1991; and year 9 extends through 1995. There are assumed to be no losses after year 9.

property since both risk levels and ultimate compensation arrangements are harder to judge.¹⁶

We know of no estimates of the aggregate cost of common risks for insurers, but data are available for one such risk, real medical malpractice premiums from 1968 to 1993 (see fig. 6.3). Medical malpractice represents a small share of total property-casualty coverage—about \$5 billion in premiums annually. It is clear that malpractice premiums vary substantially over time. Between 1968 and 1976, real premiums quintupled. Over the next decade, they nearly doubled again. Premium increases were followed by periods of real premium reductions. To the extent that these extreme premium fluctuations are the result of common-risk changes, they represent an enormous problem for insurers and the firms they are insuring.

Summary

For both single-event and common risks, the probability of a cataclysm seems quite real. Indeed, both single-event and common-risk cataclysms appear to have become more likely over time, and not merely because of the growth of coverage. In the case of single-event risks, natural disasters appear to have increased in frequency and severity since the late 1980s. In the case of common risks, the past several decades have seen an explosion of liability in areas where insurers generally did not expect it, although some reforms have ameliorated the situation.

16. With personal injuries, it is often difficult to determine what event led to the injury—hence, whether there is liability—and whether compensation will depend on casualty. The breast implant case (see n. 15 above) is one example. Malpractice is a second. Studies suggest that many errors go unpunished and that many awards do not relate to negligence (Weiler et al. 1993).

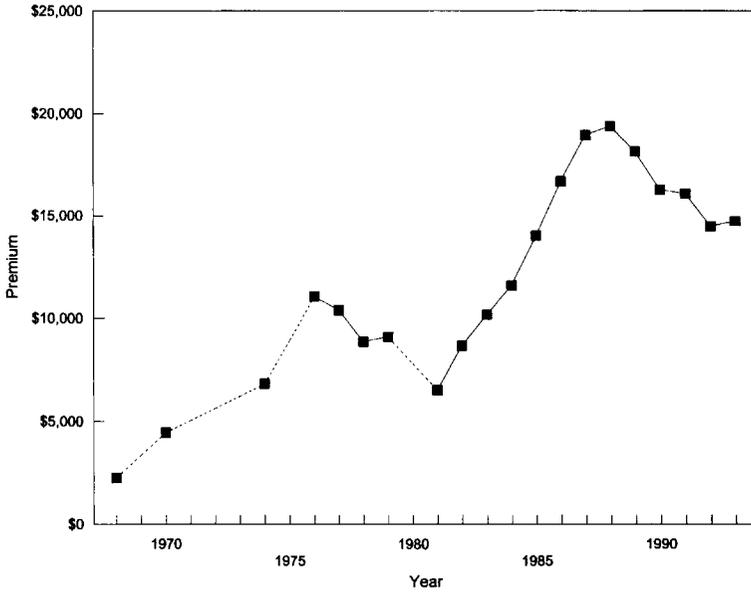


Fig. 6.3 Real malpractice premiums

Note: Data are from the American Medical Association. Dashed lines are interpolated.

6.3 Reinsurance for Catastrophes

This section investigates how reinsurance premiums for catastrophic risks should be set and policies designed in the light of potential adverse selection. Although cast in reinsurance terms, our analysis applies immediately to any situation of adverse selection and insurance, including the classic one-stage insurance transaction between an insurer and a potential insured who possesses private information. Our goal in this analysis is efficiency. We seek the policy that is optimal for the insurer, given that the risk-neutral reinsurer breaks even. Our treatment allows for a single policy with a single premium.¹⁷

We formalize the relationship between the primary insurer (buyer) and the reinsurer (seller) by assuming that there is a common prior probability distribution on the likelihood of an occurrence but that the primary insurer has special knowledge about losses since he underwrote the original risk and has a more intimate relationship with the insured. We summarize this as a signal that the buyer of reinsurance receives about the likely distribution of losses. Obviously, the buyer is more likely to seek reinsurance protection—what is often termed *ceding insurance*—when he suspects that either the probability of the insured event is high or the magnitude of a given loss occurrence is high. Such behavior is the source of the adverse selection.

17. Thus, we avoid Rothschild-Stiglitz (1976) types of difficulties—e.g., do all policies need to break even?—when there are multiple contracts in a world with multiple types.

Before turning to our second-best solutions, we should observe that a fully optimal reinsurance contract could be drawn if payoffs were made contingent on objectively observable and verifiable conditions, such as the magnitude and path of a hurricane. Beyond any problems agreeing *ex post* on what occurred, such arrangements would suffer from “basis risk”—the risk that a particular insurer’s losses would differ from average losses and thus that the insurer could not receive full reinsurance. If basis risk is significant or adverse happenings difficult to verify, we would expect reinsurance arrangements to depend on the experience of particular companies.¹⁸

There are many strategies to ameliorate adverse selection. The contractual solution is to have the insurer provide the reinsurer with detailed and accurate statistical information on the insured risks. But more will be required if asymmetries remain, for example, about the tenacity with which a claim will be pursued. One approach is to arrange in advance (before the buyer receives his private information) for the buyer to pay the seller should a risk not be sold. Such a payment would make the buyer more willing to sell, thereby reducing adverse selection. More generally, reinsuring risks shortly after they are written makes it less likely that the buyer of reinsurance has superior information.

A second strategy is nonlinear pricing depending on the proportion of the portfolio that is sold. The larger the portion of the buyer’s portfolio that he is reinsuring, the less opportunity for him to buy coverage for his worst prospects, what we term *lemon shedding* to contrast with the standard *cherry picking*. Thus, adverse selection would be discouraged if we simply reduce the per unit charge for reinsurance as a greater percentage of the portfolio is sold.

A third approach is to use proportional reinsurance to ensure that the bad risks are not the only risks reinsured.¹⁹ But adverse selection will remain a problem: insurance companies that wrote bad risks in general will be more eager to buy reinsurance than insurance companies that underwrote solid risks.

Asymmetries may cut in the opposite direction. That is, reinsurers with broader experience may be able to assess some information more precisely than are the insurers for whom they are providing coverage. This suggests that they may selectively accept policies. Whatever reinsurers may know, private information will remain on the buy side and adverse selection will persist as a problem. Moreover, if reinsurers as a group know information not known to insurers, a competitive market would prevent them from exploiting this differential knowledge. The model that follows considers only adverse selection on the buy side.²⁰

18. We are grateful to Louis Kaplow for highlighting this point.

19. Proportional reinsurance has the further advantage that, if the insured purchases multiple contracts with the same proportion, possibly from different parties, his insurance will address his total risk, which is his major concern. This is not true with nonlinear forms of reinsurance.

20. In ongoing work, we are exploring the problem of bilateral adverse selection in health care markets.

We examine nonlinear pricing mechanisms as a means to ameliorate adverse selection.²¹ We start with the observation that the losses that an insurer realizes for insured asset i will be $L_i = p_i \cdot m_i$, where p_i is the probability of a loss, and m_i is the magnitude of the loss if there is one. Both p_i and m_i may vary in the population. We denote the prior distribution of p as $f(p)$ and the prior distribution of m as $g(m)$. The means of the two distributions are \bar{p} and \bar{m} . For simplicity, we assume that p and m are independent risks, with the result that the expected loss is $E[L] = \bar{p} \cdot \bar{m}$.

Our mechanisms recognize the possibility that asymmetries may be greater on one of these components than on the other. If differential knowledge is greater on p than it is on m , or vice versa, a significantly nonlinear scheme of risk sharing is required to achieve a second-best solution. We do not expect to observe either pure case in the real world, but we do expect some cases to be characterized by asymmetry on the probability of a loss and others to be characterized by asymmetry on the magnitude of a loss. Unfortunately, the nature of the schemes in the two pure cases are quite different. This suggests that, when the buyer of reinsurance has superior information about both p and m , the second-best optimum may be relatively poor.

We develop our results by employing a simple model with a risk-averse buyer of reinsurance and a risk-neutral seller (the buyer is often said to “cede insurance” to the seller). The buyer’s utility function is $u(w) = -e^{-aw}$, where w is wealth. The seller has utility function $v(w) = w$. If we think of sellers as the reinsurance market as a whole, risk neutrality may be a reasonable assumption provided that losses are below cataclysmic levels.

6.3.1 Asymmetry on Probabilities

Consider the situation where both buyer and seller have the same subjective distribution of the size of a loss should there be a loss, $g(m)$. The buyer and the seller differ in their knowledge of p ; the seller knows only the prior probability distribution on p , $f(p)$, while the buyer has more specialized knowledge of the loss probability. For example, the seller of reinsurance may know the flood risk in a particular city, but the buyer may know it for particular houses in the city. The insurer will choose to purchase reinsurance if his houses at risk are in particularly high-risk areas. We denote the buyer’s posterior probability of a loss as p_b .²²

Since the buyer is risk averse, he is eager to lay off his loss, even at actuari-

21. Essentially, we are engaged in a process of mechanism design for a situation of asymmetric information. Our risk-sharing arrangements in effect invoke the revelation principle: the optimal second-best mechanism can always be one that makes truth telling (in this case, by the insurer) optimal. To do this, the reinsurer must commit in advance how it will use the information it secures. We are grateful to John Cochrane for reminding us of this equivalence. For more discussion, see Myerson (1979).

22. Since the asymmetry relates to a probability, it is only the mean of the buyer’s posterior distribution that matters.

ally unfair rates. Since the size of the loss is a random variable, the seller can deter the buyer from putting unfavorable risks up for sale by making him bear a disproportionate share of small losses.

The buyer has initial wealth w , which includes the premium that he has already collected. In the absence of reinsurance, his terminal wealth will be $w - m$ if there is a loss and w if there is not one. His expected utility, assuming no reinsurance, can be written as

$$(1) \quad E[U] = p_b \int -e^{-a(w-m)} g(m) dm + (1 - p_b) - e^{-aw}.$$

Ceding the Entire Risk

Since the buyer is risk averse and the seller is risk neutral, it might seem that the buyer should purchase reinsurance for the entire risk at some preestablished price. However, when p_b is low enough, the primary insurer would rather bear the risk than pay the premium for reinsurance.

Represent the price of reinsurance as r . If the payment were the ex ante actuarial value of the loss, $r = \bar{p} \cdot \bar{m}$. The buyer would purchase for values of p_b such that expected utility of being uninsured is below the expected utility of purchasing insurance, $-e^{-a(w-r)}$.

This policy may not be feasible in the market, as a simple calculation reveals. Let us say that $a = 1$, $g(m)$ is uniform on $[0, 1]$, and $f(p)$ is uniform on $[0, .1]$, implying $\bar{m} = .5$, $\bar{p} = .05$, and the ex ante actuarial premium would be $.025 (.5 \times .05)$. The buyer would purchase reinsurance whenever $p_b \geq .0352$. But, for this value, the expected loss is $.0388$,²³ which significantly exceeds the payment. No seller would take on the policy at the price of $.025$.

As the premium is increased, the buyer will be more reluctant to sell his risks. The dashed curve in figure 6.4 shows the value p_b above which the buyer will reinsure the risk, as a function of the premium for reinsurance.²⁴ The solid line is the actuarially fair premium as a function of the buyer's cutoff probability.²⁵ The point where the curves intersect is an equilibrium—the seller breaks even given the degree of adverse selection and the buyer gains from the reduction in risk.

Table 6.3 shows information about possible equilibria for various insurance schemes for the parameters described above. For comparison, the first column shows the equilibrium with no reinsurance, and the last column shows the equilibrium with perfect reinsurance. The second column shows the equilibrium when the entire risk is ceded (whole reinsurance). The risk is assumed by the

23. Since p is distributed uniformly, $E[p | p \geq p_b] = .0676$. The expected loss is therefore $.0676 \times .5 = .0338$.

24. The curve is slightly convex to the origin; a straightedge can assist the naked eye to demonstrate this.

25. We assume that the premium for the reinsurance, conditional on its being purchased, is established before the seller learns his p_b .

Copyright © 1999, University of Chicago Press. All rights reserved.

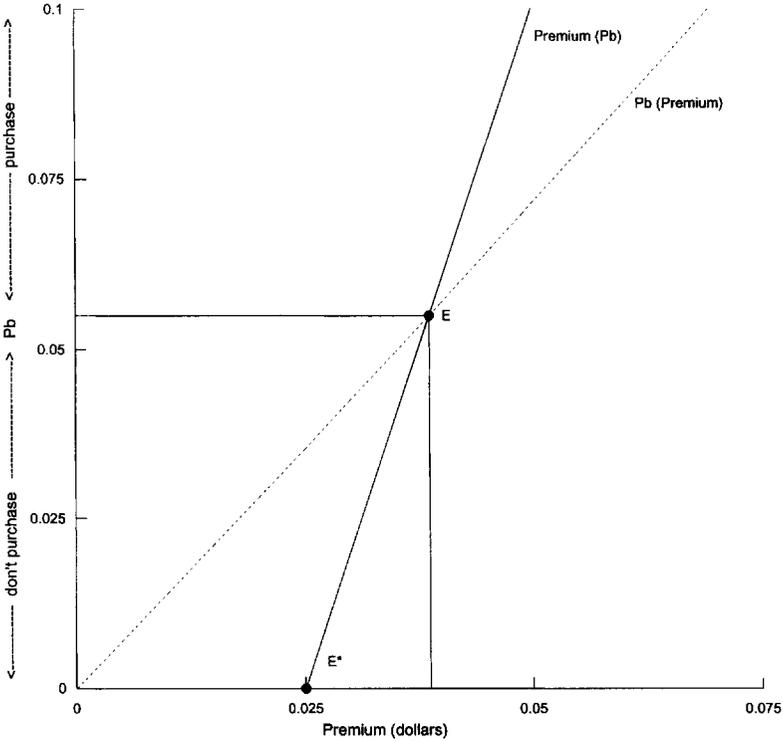


Fig. 6.4 Equilibrium in reinsurance market

reinsurer 45 percent of the time, and, because of adverse selection, the equilibrium price (.039) is well above the expected loss in the population as a whole (.025).

Whole reinsurance is quite inefficient. For values of p below p' , which occurs 55 percent of the time, there is no risk spreading even though risk spreading is socially desirable. The second row shows the percentage gain in the primary insurer's certainty equivalent, where 100 percent is the gain from perfect reinsurance. Whole reinsurance provides only two-thirds of the benefit of perfect reinsurance.

Risk aversion plays a critical role in determining when reinsurance will be purchased. Increased (diminished) risk aversion—an increase in a —tilts down (tilts up) the dashed curve and thus increases (decreases) the range over which reinsurance is purchased. If we double the degree of risk aversion locally ($a = 2$), for example, reinsurance will be purchased 69 percent of the time, and the realized gains will be 90 percent of the potential gains from perfect reinsurance.

Copyright © 1999, University of Chicago Press. All rights reserved.

Table 6.3 Comparison of Policies with Asymmetry on Probability of Loss

Measure	Form of Reinsurance				
	None	Whole	Proportion ^a	Excess of Loss ^b	Perfect
Percentage of time risk is assumed by reinsurer	0	45	50	52	100
Percentage gain in primary insurer's certainty equivalent	0	68	71	72	100
Premium (\$)039	.031	.027	.025

Note: Certainty equivalent is scaled to 0 percent in the no-reinsurance case and 100 percent in the perfect-reinsurance case.

^aReinsurer bears 81 percent of risk.

^bReinsurer bears loss above .152, where potential losses range from zero to one.

As a contrast with whole reinsurance, we investigate nonlinear reinsurance arrangements. Two such arrangements are common: proportional reinsurance and excess-of-loss reinsurance.

Proportional Reinsurance

We begin by considering proportional risk sharing. We assume that the buyer pays the seller to take on a fraction b of the risk. This fraction must be such that, with adverse selection, the premiums just cover expected payouts. For the example given above, we find that the optimal fraction is for the reinsurer to assume 81 percent of the losses ($b^* = .81$), charging a premium of .031 (see table 6.3).²⁶ The premium is closer to the population average in this case than in the whole reinsurance case because the proportional policy induces less adverse selection. The buyer's utility is greater under proportional reinsurance than under no reinsurance or whole reinsurance; 71 percent of the potential gains to the buyer are realized. The superiority of proportional reinsurance to whole reinsurance is to be expected since the latter is a special case of proportional reinsurance where $b = 1$.

Excess-of-Loss Reinsurance

Risk-averse buyers particularly want reinsurance for large losses, while risk-neutral sellers are indifferent to the size of the losses that they reinsure. This fact can be used to limit adverse selection. By reinsuring only large losses, the most important risk spreading can be achieved, with adverse selection held to a minimum.

26. The optimal policy, as with those considered later, is found by computer simulation. Some of our differences in performance across policies are small. We did not engage in the modeler's equivalent of data mining: reworking parameters and functional forms to produce bigger differences. The specific gains associated with different policies will depend on functional forms; the ranking of policies will not.

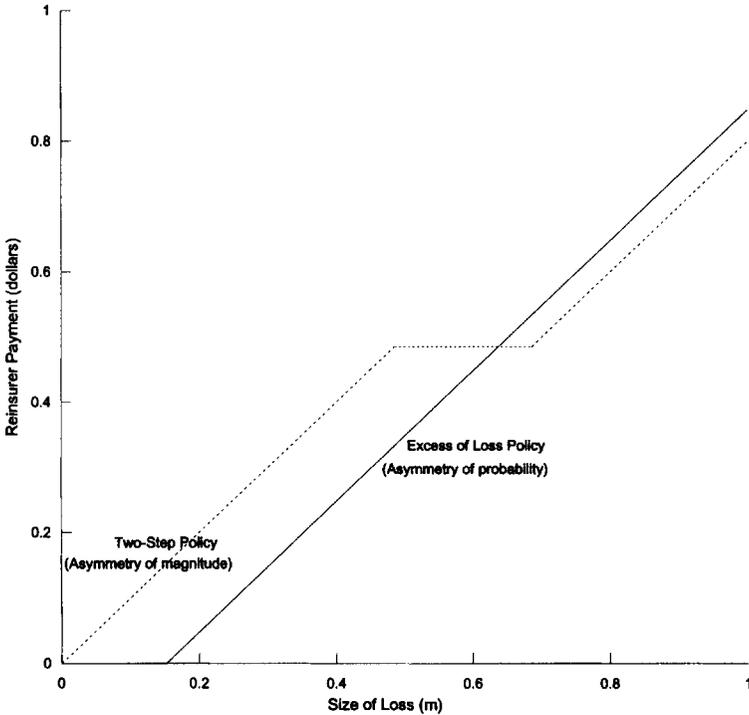


Fig. 6.5 Optimal reinsurance with adverse selection

We consider arrangements for nonproportional, or excess-of-loss, reinsurance. As noted above, this is a common form of reinsurance in current markets. We constrain such policies (and all policies in this paper) in two ways: the buyer’s share of the loss can never exceed 100 percent; and, at the margin, the buyer’s share of losses can never exceed 100 percent.

The optimal reinsurance arrangement subject to these constraints is shown in figure 6.5 and described in table 6.3. The policy has the buyer of reinsurance assume all risks up to $m = .152$. Beyond this point, the reinsurer assumes all losses (hence the title *excess-of-loss*). Premiums are only slightly above the no-selection level, and 72 percent of the potential gains from reinsurance are realized. Indeed, we can demonstrate that, when there is asymmetry about the probability of a loss, excess-of-loss reinsurance is optimal.²⁷

PROPOSITION 1. When the buyer has no special knowledge of m but knows more than the seller about the probability of a loss, the optimal reinsurance

27. An excess-of-loss policy will also be optimal when there is moral hazard on the probability of a loss (i.e., having insurance affects the loss probability) but not on the magnitude of a loss. In ongoing work, we are exploring optimal reinsurance with moral hazard.

Copyright © 1999, University of Chicago Press. All rights reserved.

policy charges the risk-averse buyer for all losses up to a point, with the risk-neutral seller responsible for losses beyond that point.

PROOF. See the appendix.

Given the theoretical merits of excess-of-loss reinsurance and its common occurrence in the real world, it might be logical to conclude that it is always the most-preferred arrangement. This would be a startlingly inappropriate conclusion.

6.3.2 Asymmetry on Size of Loss

We now consider situations where the primary asymmetry is on the magnitude of loss, not its probability. For example, the insurer understanding the litigation climate for toxic torts may have expectations about punitive damages that are not available to more distant reinsurers. Asymmetry about size of loss may apply as well for natural disasters. For example, a Korean insurance company will likely have superior information relative to an overseas reinsurance about whether its clients' buildings would survive an earthquake, whether fire services would be sufficient should there be one, etc.

To simplify discussion, we assume that both the buyer and the seller of reinsurance believe in common that the probability of a loss is p . However, assessments of damage conditional on loss are asymmetrical. Specifically, we assume that the buyer knows the damage from a loss exactly, m_b ,²⁸ and that the reinsurance seller has loss distribution $g(m)$. Thus, the buyer knows everything, while the seller knows nothing.

Table 6.4 shows a variety of policies, each optimized for a different structure, given this type of information asymmetry. For excess-of-loss reinsurance, it is optimal for the reinsurer to start to pay at zero. In effect, the optimal excess-of-loss plan offers whole insurance. This in turn makes it more attractive to buy insurance when losses will be small, helping deter adverse selection. If there were a threshold for reinsurance coverage, the buyer would not purchase reinsurance for any losses below the cutoff, indeed, for some amount above it. Adverse selection would be severe. Even with the optimal plan of this type, risks are not laid off 60 percent of the time.

Interestingly, when information asymmetry is about the magnitude of a loss, proportional reinsurance is superior to excess-of-loss reinsurance. The optimal policy has the reinsurer assume 89 percent of the risk, and the policy is purchased 42 percent of the time. Even with proportional reinsurance, however, adverse selection has a large effect, as indicated by the substantial difference between premiums and average losses.

28. In the case of asymmetry about probabilities, a sufficient statistic for the buyer's superior knowledge is the mean of his distribution, p_b . That is because rational decision makers are not risk averse on uncertainties about probabilities. With uncertainty about expected losses, we would need to summarize the entire probability distributions of the buyer and seller to characterize the market accurately. For simplicity, we assume that the buyer knows the actual value of a loss should it occur.

Table 6.4 Comparison of Policies with Asymmetry on Magnitude of Loss

Measure	Form of Reinsurance					Perfect
	None	Whole/ Excess of Loss ^a	Proportion ^b	Whole with Cap ^c	Two Step ^d	
Percentage of time risk is assumed by reinsurer	0	40	42	49	54	100
Percentage gain in primary insurer's certainty equivalent	0	80	81	85	87	100
Premium (\$)040	.035	.037	.029	.025

Note: Certainty equivalent is scaled to 0 percent in the no-reinsurance case and 100 percent in the perfect-reinsurance case.

^aOptimal excess-of-loss reinsurance is whole reinsurance.

^bShare of time risk borne by reinsurer is 89 percent.

^cPrimary insurer starts paying at .686.

^dPrimary insurer starts paying .486 and stops paying at .686, where potential losses range from zero to one.

The next two columns show alternative reinsurance designs. The fourth column examines a “whole-with-cap” policy—a policy that has the reinsurer fully assume small losses and the primary insurer pay for all losses above a cutoff (the optimal cutoff is .686). This is the opposite of an excess-of-loss policy. The whole-with-cap policy is superior to the excess-of-loss policy; it is purchased 49 percent of the time, and 85 percent of the potential utility gains are realized.

The last column considers a reinsurance arrangement where the reinsurer bears all small losses, the primary insurer bears a layer of losses, and the reinsurer then reassumes the risk. We term this a *two-step* policy.²⁹ This policy outperforms all the other policies examined; it is purchased 54 percent of the time, and 87 percent of the potential utility gains are realized. Our simulations suggest that the two-step policy is the optimal reinsurance arrangement, although we do not currently have a formal proof that it is always best.

Figure 6.5 compares the reimbursement structure of the optimal two-step policy and the optimal excess-of-loss policy from table 6.4. The optimal policies for these two cases are strikingly different. While these results may appear puzzling, the intuition is readily grasped. Optimal reinsurance makes the buyer bear more of the risk on the dimension about which he has superior knowledge. If the buyer knows more about the probability of a loss than its magnitude, reinsurance should exclude coverage for the common small risks that are the source of exploitable adverse selection; this is exactly the “deductible-style”

29. The optimal two-step policy has the first layer of no reinsurance up to losses of .486, full reinsurance up to losses of .686, and then no reinsurance above that point.

Copyright © 1999, University of Chicago Press. All rights reserved.

policy that characterizes excess-of-loss insurance. If the buyer knows more about the magnitude of loss than its probability, however, it is optimal to make it costly for him to buy reinsurance when he knows that the magnitude of loss is high. This is why the capped and two-step policies are desirable when the asymmetry is on the size of loss.

6.3.3 Summary

Optimal reinsurance policies differ significantly depending on the form of information asymmetry. If the asymmetry in information is solely about the probability of a loss, the traditional excess-of-loss reinsurance contract is optimal. If the asymmetry is only about the size of a loss, however, it is preferable for the buyer to assume extreme losses as well as ordinary losses. Most real-world situations will involve elements of both asymmetries; hence, they will require future analysis to define the optimal policy structure.

Current reinsurance arrangements predominantly use proportional or excess-of-loss reinsurance. If most of the asymmetry between buyers and sellers is about the probability of a loss, our results suggest that this is reasonable. However, if asymmetry is primarily about the magnitude of a loss rather than its probability—as seems reasonable for many catastrophes—these types of insurance arrangements may be far from optimal. Thus, reinsurance arrangements might be considerably improved. Indeed, since our analysis carries through to markets for primary insurance coverage as well, such contracts also might be improved.

6.4 Reinsurance for Cataclysms

Many important risks are common to all insurers. For example, an earthquake in California will affect all insurers that provide coverage in that area. Similarly, changes in automobile-accident rates or tort-liability judgments for accidents will affect all insurers underwriting automobile insurance. Reinsurance can help diversify losses for any one insurance company, but, if the aggregate loss is sufficiently large, the capital of the worldwide insurance and reinsurance industries may not be large enough to cover the losses.

Cataclysms resulting from some common-risk changes are already spread relatively widely. If automobile-liability costs in the United States go up by 10 percent in a year, for example, that may be serious for the automobile-insurance industry, but the risk will be widely spread since the geographic basis is wide and many different insurers will be involved. Far more significant would be an event with the same total damage but in a geographic region where some insurers are concentrated, such as an earthquake. Of course, diversification across areas of the country is less effective when a common risk happens and losses escalate.

Both single-event and common-risk cataclysms have far-reaching implica-

tions. The Insurance Services Office, for example, estimates that a cataclysm on the order of \$50–\$100 billion could bankrupt one-third of all insurers. Unreimbursed citizens will lose, and the government will almost certainly incur high costs in reconstructing public facilities, offering disaster aid to uncompensated victims, and suffering losses in tax revenues.

6.4.1 Market-Based Responses to Cataclysmic Risk

The market appears to have perceived the increase in cataclysmic risk in recent years and has taken steps to limit coverage for it. Often, this is accomplished by withdrawing from the market. In the past few years, for example, insurers and reinsurers have withdrawn from homeowner's insurance in Florida (as Nationwide did recently) and California. Insurers and reinsurers also dropped coverage for terrorism risk in the United Kingdom after bombings in the early 1990s. These actions can make it impossible for people to buy insurance even for noncataclysmic losses. When insurers withdraw from the Florida homeowner's market, for example, homeowners cannot insure themselves against even relatively minor losses.

With many common-risk problems, insurers have responded by effectively shortening the life of the policies they sell. Medical malpractice and environmental liability are two common examples. In both these markets, traditional insurance was occurrence based. As the nature of long-term risk became more apparent in the 1970s and 1980s, insurers switched to a claims-made insurance policy, one that covers the policyholder only for claims filed in a particular period of time for damages occurring during that time.

Insurers may prefer such mechanisms to raising their premiums to cover cataclysmic risk. Political or social considerations, for example, may prevent insurers from increasing premiums to the level required to insure against increasingly important cataclysms. In addition, the managers of the insurance company may not want to risk insolvency even though the expected profits may be great.³⁰

Table 6.2 above shows some evidence of how the market's response to growing cataclysmic risk affects firms and individuals seeking insurance. Under an occurrence-based malpractice-insurance policy, only 6 percent of losses are realized in the first year, and only two-thirds are realized within five years. In contrast, under a claims-made malpractice policy, 20 percent of the losses are borne in the first year, and 85 percent of the losses are borne within five years.³¹ Hence, a claims-made policy has less long-tailed risk. Ignorance or uncertainty

30. In theory, shareholders might like charging actuarial prices for risks that would not pay off in case of occurrence. Their losses on the downside are capped, raising the expected value of their investment.

31. This ignores any differences in the types of policies that are sold on a claims-made or an occurrence-based basis. Because the longer-risk types of medical care are more likely to be sold as claims-made policies, this should reduce the loss difference in the two policies.

about the losses incurred will be reduced, and premiums can be set on a sounder basis. Claims-made policies have had a similar effect on product-liability coverage, although not other liability coverage. Curtailing the life span of insurance coverage returns the risk to insureds, whose premiums will rise and fall with estimates of coverage costs. This risk falls on the insured in part because most of the events covered by new-insurance purchases will have already happened but not yet been compensated.

The potential losses from common-risk cataclysms thus seem substantial. The market's response to date has been largely to increase premiums for insurance, to reduce the extent of coverage, and to limit exposure to long-tailed losses. The net result of these actions is to put much more risk back on the purchasers or insurance. The risk of premium variability over time thus becomes a concern. The rest of this section examines other actions that could be taken to increase financial protection against cataclysmic losses.

6.4.2 Reserves and Borrowing Capacity

Typically, insurers build reserves so that they can pay for very large damages. However, reserves are unlikely to be an effective means to deal with cataclysms.

It is a challenge for mutual insurance companies to build up substantial reserves. Current policyholders, who may not be insured in the future, are likely to think that the reserves are their monies. Saving them for future policyholders who may suffer cataclysms is hardly in their interest. Beyond this, when reserves are significant, there is tremendous pressure on companies to demutualize and thereby distribute them.³²

For-profit insurers or reinsurers could build reserves because the money comes from and belongs to their shareholders. But the primary reason to build up reserves would be to reassure policyholders that the company will survive the cataclysm. Given the difficulty of communicating information on the soundness of one's company to policyholders, we would expect for-profit companies to have inadequate reserves for outlier losses.

The government could regulate reserve levels to be adequate to meet a cataclysm, but it might be a quixotic quest. If we wish an insurance company today to build up reserves against losses that are ten times an ordinary year's losses—as, for example, a major hurricane might impose—then it would have to charge an extra premium of 50 percent a year for twenty years to have enough money on hand. Not much is accomplished if the event occurs in year 7. And, while a regulator could require a for-profit reinsurer to raise enough capital to be prepared for a cataclysm, the reserves required may simply be too large to be practicable.

32. Demutualization has been widely observed for savings banks and could be an increasing phenomenon for insurance companies. Some mutuals have large reserves.

An alternative approach to protecting against cataclysms is to have deep-pocket parties post bonds or their equivalent. This was the approach of Lloyd's of London. Rich people essentially placed their whole net worth at risk to cover potential future claims. The great advantage of this "pledge" system for reserves is that monies are kept working in high-valued uses until needed. The system proved less effective when it was called on severely, as Lloyd's was when the rapid escalation of costs for asbestos and other changing liabilities exhausted its borrowing capacity. Even people who have made large amounts of money over many years, as many of the Lloyd's names did, get upset when a substantial portion of their net worth is lost. Lloyd's is now reconstituting itself, but with more restricted obligations of outside investors; presumably, this has diminished the security of its reinsurance capacity.

Outside Lloyd's-type pledge arrangements, it may be possible for reinsurers to borrow to pay out cataclysmic losses, with the promise of future premiums as collateral. However, such borrowing might not be repaid for up to a century or more (in the case of a major earthquake) and would presumably be offered only at a substantial risk premium.

In sum, traditional methods of diversification are unlikely to work well in the case of cataclysms. The next subsections discuss two alternatives that might be more effective.

6.4.3 Securitization of Cataclysmic Risk

Cataclysms may be large for the reinsurance industry but are generally small for the world's capital markets as a whole. A \$100 billion earthquake loss, for example, would bankrupt a fair share of the reinsurance industry but would represent only 2 percent of the New York Stock Exchange asset value and less than 1 percent of the value of the U.S. capital market. Movements of this magnitude in net asset value in a month are the norm. Further, many cataclysms will be uncorrelated with economic conditions, and investments in cataclysmic risk could therefore help market investors diversify their portfolios.

Thus, it seems promising to diversify cataclysmic risk in broader securities markets, much as is now done for home mortgages. Such diversification for risks is beginning to occur.³³ Perhaps the best-known examples of the securitization of sizable risks are the catastrophe options that have been traded on the Chicago Board of Trade since 1992. For example, reinsurers can buy a call option that pays off if aggregate losses are above a given level—perhaps \$5 billion. The size of the payment would be directly proportional to the losses above \$5 billion. The call could be sold by firms that would do well in the event of a catastrophe—home builders, for example—or by general investors looking for a chance to diversify outside traditional securities markets.

33. For a discussion of securitization of cataclysmic risk, see Borden and Sarkar (1996) and Litzenger, Beaglehole, and Reynolds (1996).

Other instruments for trading cataclysmic risks have been introduced, but they have been less popular than options. Contingent surplus notes, for example, guarantee the insurer buyers for its debt in the event of a cataclysmic loss. Typically, money is taken up front and invested in Treasury securities, which can be exchanged for company bonds at the insurer's discretion if a cataclysm occurs. Nationwide Mutual completed such a transaction recently. Investors purchased \$400 million in U.S. Treasury bonds, and Nationwide can convert the Treasury bonds to company notes in the event of a catastrophe. Thus, Nationwide has access to a ready source of cash in the event of a catastrophe. Investors are compensated by getting a higher return than a normal Treasury bond would yield.

Catastrophe bonds (also termed *act-of-God bonds*) also yield cash from catastrophes for insurers. These bonds have coupon payments that depend on the level of the insurer's losses. As losses increase, insurers are not required to pay out as much and can use interest on the initial principle (and sometimes the principle itself) to pay out the unusually high claims.

One concern with all these financial instruments for offering reinsurance protection is whether the triggering event should be company specific or marketwide. Insurers and reinsurers would naturally want to have company-specific securitization, other factors equal. Basis risk is an important issue for insurers, and diversification against that risk would bring substantial value (Major 1996). But company-specific securities make both adverse selection and moral hazard more likely. For example, insurers that know that they have financial market reinsurance might pay less attention to the loss probabilities or magnitudes of their insureds. A likely scenario is that securitization will be based on marketwide losses and that intermediaries will pool particular companies into smaller groups approximating the market average, where risk spreading can be achieved but moral hazard and adverse selection minimized.

To date, securitization of cataclysmic risk has been limited. Catastrophe options, for example, trade less frequently than do most other options on the Chicago Board of Trade, and contingent-surplus notes and act-of-God bonds trade relatively infrequently. This relative illiquidity has probably led to a higher required return for investors than would be required if trading were more fluid. It is unclear whether the market will ultimately become large enough to support true cataclysmic diversification or whether other forms of reinsurance will be required.

6.4.4 The Government as Reinsurer

The government has a deep credit capacity: it can borrow by issuing debt far more readily than can private insurers or reinsurers, and it can raise resources rapidly through its ability to tax.³⁴ The ability to borrow or tax to meet

34. The government might also build up reserves to pay for potential cataclysms, but such build-ups are rare in the public sector, as the examples of social security and Medicare demonstrate.

cataclysmic obligations depends only on political capabilities. This puts the government in a natural position to be a reinsurer. The government can also compel reinsurance to solve the adverse-selection problems noted above. This is one of the traditional rationales for mandatory participation in social security, Medicare, and other public-insurance programs.

The government is also a natural reinsurer because it is often called in to cover cataclysmic losses after they have occurred, whether it has contracted for them or not.³⁵ In the recent past, the government has bailed out many savings-and-loan institutions; these efforts will cost taxpayers money for many years to come.³⁶ The government will pay the cost of cleaning up its atomic-weapons waste sites, estimated at well above \$100 billion. The Price-Anderson Act, which is no longer in effect, insured nuclear utilities against severe losses from individual accidents; the cutoff of private loss was \$560 million. Governments also run insolvency funds that bail out policyholders when their insurance company goes bankrupt. And, when hurricanes or floods occur, the government often provides disaster aid.

Formal government reinsurance programs are often established when primary-insurance markets break down. For example, the U.K. government decided to provide terrorism insurance following the breakdown of insurance and reinsurance markets when the bombings in London began in 1992. High-risk homeowners' pools in the United States typically get under way when the premiums that the private market would charge are above the levels that are socially acceptable.

Establishing a formal reinsurance system, and thus collecting at least some money, may be preferable to waiting until the event occurs and spreading the costs more broadly. Selling insurance in advance of a disaster also makes people aware of the social cost of their actions and thus limits moral hazard—for example, discouraging excessive building on floodplains.

The government could price reinsurance at its expected cost or, alternatively, provide a subsidy to purchasers. A subsidy might be appropriate if the government is worried that people will not purchase insurance appropriately or if the government wants to share in the losses from a catastrophe the same way it shares when it receives additional tax revenue when there are not catastrophes. Flood insurance, for example, is heavily subsidized. Savings-and-loan insurance, pension-plan insurance (PBGC), and state insurance-guarantee funds offer a more subtle subsidy: they charge premiums but typically fail to differentiate the premiums on the basis of risk. For example, savings-and-loan premiums were the same for all institutions. As we saw, this encouraged risky activity for institutions near default. Similarly, the PBGC used to provide pension insur-

35. Indeed, the probability of government involvement in certain cases may explain why insurance is only rarely purchased for some events (Kaplow 1991).

36. The government also has the authority to order other parties to pay for damages, as it did with Superfund. The government's unique capability to pass and enforce legislation—up to the limits of constitutionality—puts it in a position to get someone else to pay for removing a risk.

ance at \$19.00 per participant plus \$9.00 per \$1,000 in vested but unfunded benefits. However, there was a cap of \$53.00 per participant. This created an incentive to defund a plan once the cap was approached or passed. Many insurance-guarantee funds charge the same rate for all insurers, thus encouraging insurers to gamble with state funds if they are near bankruptcy and eliminating the need for individuals to think about the solvency of the company from which they purchase insurance.

Because of these moral hazard problems, the government might want to provide insurance at an actuarially fair cost. Indeed, if the government is providing reinsurance solely because its size allows it to better diversify cataclysmic risks, that is a function that the market should be willing to pay full cost for. Some insurance has been proposed along these lines. The Treasury Department, for example, recently proposed public reinsurance for natural disasters with aggregate losses from \$25–\$50 billion. Reinsurance would be auctioned to insurers in amounts ranging from \$1 to \$25 million. Although this proposal has somewhat of a market flavor, there are significant questions: Would the amount auctioned be too great to secure the government at least actuarial prices? Would such insurance be pitched at too low a range, given the government's inevitable financial responsibilities when losses exceed \$50 billion?

In other cases, government insurance is moving closer to actuarial pricing. For example, the Retirement Protection Act of 1994 eliminated any cap on PBGC per participant costs starting in 1997. As a result, the PBGC should come closer to risk-based pricing.

One might question how the government will know enough about appropriate pricing to sell actuarially fair insurance when the government is the only insurance seller. The recent experience of the California Earthquake Authority (CEA) does not suggest that government authorities are well equipped to secure low prices for reinsurance. The CEA was seeking to purchase reinsurance for the layer between \$7 and \$8.5 billion over a four-year period. Its best estimate (based on scientific information) was that the layer would be reached with probability 1.2 percent, implying a fair market premium of \$18 million. In the end, Warren Buffet sold the CEA this layer of reinsurance for \$161 million per year. This is roughly ten times the estimated actuarial cost.

To learn what prices the government ought to use in selling reinsurance, the government may wish to inject a sliver of private-sector pricing of cataclysmic risk to indicate how much it costs to provide. For example, the government could sell off a small portion of its book, say a total of 1 percent in 0.1 percent units. Then, if the government were responsible for \$100 billion of risk, investors could set prices for accepting risks up to \$100 million. If the market price for a 0.1 percent share turned out to be \$10 million, this amount would be paid to the investors, who would then take on the responsibility for this portion of the risk. Equally important, this rate would be charged to insurers who are reinsured by the government. Depending on administrative complexities, these

prices could be auctioned to the private sector on a basis that reflected the risks involved, the specific insurer, or both.

The issue of basis risk adds complexity to such an auction scheme. Demand for reinsurance is likely to vary with the risk exposure of the particular company. Companies with a broad nationwide exposure will find a claim based on aggregate losses well suited to its needs. Companies with concentrated risks in a particular area or type of catastrophe, however, will find such reinsurance policies less suitable. This variation in the ability to obtain hedging suggests that the price of the reinsurance policy may vary with total sales. The equilibrium price will be high if little reinsurance is offered but low if a lot of reinsurance is offered. Determining how much reinsurance the government should provide in such a situation is a difficult issue.

Many countries use public insurance to address cataclysmic risk (although, to our knowledge, the private sector has not been involved in setting prices). Table 6.5 shows examples of public-sector reinsurance programs for catastrophic risk: France, the Netherlands, Japan, Spain, New Zealand, Norway, South Africa, and the United Kingdom. Most of the reinsurance programs cover natural disasters. In France, for example, flood losses that are particularly large are paid for by the national reinsurance company. The company collects premiums from a mandatory assessment on insurance premiums for property damage, business interruption, multirisk policies, and automobile insurance. Other countries cover events such as terrorism or political risk.

Public-reinsurance systems have three advantages beyond their ability to address cataclysmic risk. First, if the government provides reinsurance for cataclysms, it might allow a private market to develop for smaller aggregate losses. Second, public-sector reinsurance may address a moral hazard problem concerning government behavior. Many cataclysms can be ameliorated or magnified by government policy. Legal judgments, for example, vary with legislative and judicial actions. Decisions about where houses can be built will affect hurricane and earthquake losses. By making the government the reinsurer for

Table 6.5 Public-Sector Catastrophic Reinsurance Programs

Country	Year Began	Coverage
New Zealand	1944	Earthquake; war damage
Spain	1954	Natural disasters; terrorism; peacetime acts of armed forces
Japan	1966	Natural disasters
South Africa	1979	Political riots
Norway	1980	Natural disasters
France	1982	Natural disasters
The United Kingdom	1992	Terrorism
The Netherlands	1996	Natural disasters

these risks, the government may internalize some of the costs of its actions. Third, if government is the insurer of last resort anyway, a formal public-reinsurance system may raise revenues for this service. Of course, if insurance would be provided without the government, government involvement may just lead to a subsidy of risky activities that is not warranted.

6.4.5 Cataclysms and Learning about the Future

The occurrence of a cataclysm may affect our predictions about the magnitude of losses in the future. On one side, a cataclysm may eliminate the highest-risk situations; an earthquake may bring down the weakest buildings and release stress in a fault. Moreover, once society is alerted to the risk, it may take actions, say, reinforcing buildings, that will make the risk less likely or less costly in the future.

Cataclysms may also portend greater future losses if they result in an updating of event probabilities or the distribution of loss magnitudes. If we think the weather may be changing, then a hurricane this year may suggest additional hurricanes in future years. And, if losses are surprisingly high given a hurricane, we may conclude that future hurricane losses will be greater than we had anticipated.³⁷ This is anecdotally what occurred after Hurricane Andrew in Florida.

When cataclysms lead to updating in this fashion, a second major risk arises—the risk that premiums will rise in the future (Cochrane 1995; Cutler 1996). For example, the liability crisis of the 1970s and 1980s was caused much more by a rapid escalation in premiums than by any set of events that in themselves imposed substantial costs.³⁸ Just as costs escalated dramatically for insureds, so too did they for insurance companies seeking to reinsure risks. Insurers, presumably being more risk averse than those selling reinsurance, would probably have liked to have insured their future premiums as well.³⁹ But this insurance was not provided.

Indeed, reinsurance to protect against future premium increases is rarely issued, for three reasons. First, the insured has a quantity choice to make (the amount of insurance to purchase), which would be distorted if rates were not equal to current actuarial costs. Second, no reinsurer could bear the risk associated with protecting its customers against future increases in premiums. It would be offering protection against common-risk cataclysms. The exposure would be monumental. Third, a large up-front payment would be required to

37. For most statistical models, the more extreme the event, the more significant is the updating of future probabilities. Thus, the occurrence of a \$25 billion hurricane will increase estimates of future event likelihoods and costs more than the occurrence of a \$5 billion hurricane.

38. Macroeconomic factors such as inflation and high interest rates were also important in the liability crisis.

39. This is analogous to having future health-insurance premiums guaranteed even if the person contracts a chronic, expensive disease.

prevent insurance purchasers from leaving the contract if insurance premiums decline. Charging in this fashion seems impractical.

Current insurance arrangements generally put the buyer of insurance at risk for premium increases and decreases over the period of the policy, usually one year or (in a claims-made policy) somewhat longer. If there is little movement in expected costs from year to year, say, because there is little learning, this risk is minimal. However, experience over the past several decades suggests that expectations of future costs change substantially, and thus variability in premiums is large.

To determine how important learning about future losses is, we use the PCS data to determine how predictions of losses should be updated over time. We assume that expected losses in any year are a function of expected losses in the previous year and the actual loss experience in the previous year:

$$(2) \quad E[L_{t+1}] = \alpha \cdot L_t + (1 - \alpha) \cdot E[L_t].$$

The coefficient α is the weight on the previous year's losses.⁴⁰ Therefore, α will be greater than zero if the world changes fundamentally over time in a way that insurers had not predicted. For example, if the underlying cause of North American hurricanes (perhaps weather conditions in sections of Africa) changes in such a way that more hurricanes are likely, this would be represented first as an increase in current hurricane losses, which then feeds back into expected future hurricane damage.

Actual damages in any year are expected damages times an error term:

$$(3) \quad L_{t+1} = E[L_{t+1}] \cdot e^{\varepsilon_{t+1} - .5\sigma^2},$$

where ε is normally distributed with variance σ^2 . The second term in the error term is an adjustment to make the mean of actual losses equal to the mean of expected losses.

Equations (2) and (3) can be estimated for α and σ^2 and expected losses in the first year of the data. Combining all the catastrophes, our estimate of α is .02, with a standard error of .01.⁴¹

To understand what this implies about premium variability, we simulate premiums and losses using equations (2) and (3). We simulate thirty thousand separate three-hundred-year time-series paths. We exclude the first one hundred years of the simulation to avoid start-up effects, with the result that we have effectively 6 million years of data.

Table 6.6 summarizes these simulations. The first columns show results for $\alpha = .02$; the second column shows results for $\alpha = .20$. The higher level of α

40. Estimates using average losses in the previous few years instead of just the past year yielded very similar results.

41. We attempted to estimate the model separately for different types of risks (particularly hurricanes) but found the estimates of α imprecise for risks with large variability.

Table 6.6 Simulated Losses from Premium Variability (\$ million)

Loss Distribution	Weight Placed on Current Year's Losses			
	$\alpha = .02$		$\alpha = .20$	
	Premium	Actual	Premium	Actual
Overall distribution:				
Mean	2,651	2,652	2,724	2,725
Standard deviation	876	3,273	73,433	100,258
Percentiles of average values:				
10th percentile	1,801	1,666	13	11
Median	2,509	2,471	159	144
90th percentile	3,670	3,856	2,539	2,416
99th percentile	5,230	5,756	33,326	33,231
Average of percentiles:				
Mean 10th percentile	2,256	559	392	188
Mean 90th percentile	3,084	5,739	7,127	7,049

Note: Based on 30,000 simulations each of 200 years' duration. The degree to which expected losses change with current-year losses is given by α . The second block shows the distribution of the average premium or loss across the 30,000 simulations. The third block shows the distribution of the 10th percentile premium or loss within each simulation and the 90th percentile premium or loss within each simulation.

is meant to proxy for a market with a greater degree of persistence in innovations. The first rows of the table show the mean premiums and losses across all years of the simulation. Not surprisingly, mean premiums match mean losses for each value of α .⁴²

Our first measure of premium variability is the distribution of the average premium. That is, if an insurer cares only about the average premium over a two-hundred-year period, how variable is that average premium? This would be the relevant measure of variability for an insurer that was able to borrow and lend costlessly but was worried about what its average repayment over a long period of time would be.

We show the distribution of average premiums and losses for each two-hundred-year simulation in the second block of the table. Over a two-hundred-year period, average premiums are distributed essentially the same as average losses. For each value of α , the tenth percentile of average premiums is essentially the same as the tenth percentile of average losses, and the same is true at other points in the distribution. Thus, for a given level of α , insurance helps spread the risk that actual losses in a year will substantially exceed the mean, but it offers little benefit in reducing the spread in average losses for the two-hundred-year period. There are large differences in the variability of premiums as α changes, however. The ninety-ninth percentile average premium for the

42. The mean loss for $\alpha = .20$ is above the mean loss for $\alpha = .02$ because of sampling variability.

case of $\alpha = .20$ is over six times greater than the ninety-ninth percentile premium for the case of $\alpha = .02$.

Not all insurers are risk neutral over a two-hundred-year horizon, however. Bankruptcy fears or agency costs may make insurers care about losses at a greater frequency than the long-run average. Our second measure of variability captures this issue. We show the average value of premiums or losses across simulations. That is, if one thought about the ninetieth percentile premium *within* the two-hundred-year period, what is the average value of that number? To the extent that the average value of the ninetieth percentile premium is substantially above the average value of the average premium, the within-time-period risk will be substantial.

We show the distribution of these percentiles in the lower panel of the table. It is clear that insurance has an important role to play in this case. For the case of $\alpha = .02$, the average value of the ninetieth percentile premium is 45 percent below the average value of the ninetieth percentile actual loss. But the ninetieth percentile premium is 16 percent above the premium that would be required if insurers could reinsure the risk associated with variation in expected losses (the premium for this would be the overall average loss, \$2.7 billion). For the case of $\alpha = .20$, the ninetieth percentile premium is 160 percent above the unconditional mean loss. The losses to an insurer from premium variability due to changing loss processes may be substantial.

While current insurance arrangements generally either fix premiums over time or vary them fully with annual changes in expected losses, respectively placing all risk on the insurer or insured, other arrangements are possible. Some would share the burden of premium increases between the insurer and the insured. This type of “premium sharing” could perform a function much like insurance and risk sharing. The optimal form of premium sharing, of course, will depend on the utility function of the insured and insurer. If the two parties have exponential utility, albeit with different parameters, then they should bear local and global risks in proportion to those parameters. Side payments at the outset provide for movements along the Pareto frontier. By contrast, if both parties have utility functions where wealth is raised to a power, the logarithm being a special case, then they should proportionally share the total pie, including initial wealths. Movements along the Pareto frontier are achieved by varying the sharing percentage. In both situations, once a division of welfare is agreed on, all risks are shared in the same proportion (see Wilson 1965).⁴³ In our model, the risks are experienced as expected costs per period and in variability about those expected costs.

In theory, if firm contracts could be drawn for such situations, the parties would share increases and decreases in premiums in the same percentage that they shared any actual losses. In practice, it will not be possible to work out

43. Pratt and Zeckhauser (1989) show that this is the only case for which sharing is proportionate.

such sharing arrangements owing to the insecurity of premium promises. But we should recognize that the present situation, where insureds pay a small proportion of losses when period losses are great but bear virtually all the variability in premiums, is far from optimal.

6.5 Conclusion

Catastrophic and cataclysmic losses are an increasing concern for insurance companies, reinsurers, and society. In recent years, losses have escalated because of both unpleasant surprises in natural disaster magnitudes and a changing liability game. Some major insurers have become insolvent; many have been restructured or sold under duress. Insurance difficulties impinge on the public through premium boosts and the disappearance of insurance in some markets. Ultimately, economic activity suffers.

Such developments suggest that it is important to insure the insurers. Reinsurance markets are booming, but past reinsurance structures may be inadequate to today's markets. Reinsurers are likely to be less familiar with the risks that they are insuring than their insureds, both because the reinsurance market is increasingly globalized and because new risks, such as those related to technology or terrorism, need to be insured and reinsured. As a result, asymmetrical information and adverse selection are likely to be problems. If the asymmetry is on the magnitude of the loss, standard reinsurance frameworks—such as excess of loss or proportional—are inappropriate.

Cataclysmic losses—those that overwhelm insurance markets—are more common than they used to be. Cataclysms can stem from both single events and common risks. While new financial instruments are bringing additional liquidity to insurance markets, the government has substantial advantages as a reinsurer given its deep credit capacity. Although putting government insurance on a sound financial basis, say, by taking guidance from private transactions, is perhaps an unlikely step, it would contain the dangers of bureaucratic optimism and political interference.

Nature and the economy, not to mention government processes, are throwing risks at us on a grand scale. Significant innovation will be required if insurance and reinsurance markets are to make a sufficiently grand response.

Appendix

Proof of Proposition 1

The insurer can choose to purchase reinsurance or not. Denote utility if the insurer does not purchase reinsurance as U_N and utility if the insurer purchases reinsurance as U_I . These utility functions are given by

$$(A1) \quad \begin{aligned} U^N &= (1 - p_b)U(Y) + p_b E_m[U(Y - m)], \\ U^I &= (1 - p_b)U(Y - \pi) + p_b E_m[U(Y - \pi - \theta(m))], \end{aligned}$$

where Y is premium income, π is the reinsurance premium, p_b is the probability (known to the insurer) of experiencing a loss, m is the realized loss, and $\theta(m)$ is the insurer's payment if he has purchased reinsurance. E_m signifies that the expectation is taken with respect to the random variable m . The first term in each utility function is utility if there is no loss. The second term is utility when there is a loss.

Let p_b^* be the probability where the insurer is indifferent between purchasing and not purchasing reinsurance ($U^N = U^I$). We can express the difference in utility $U^I - U^N$ as the sum of two terms: the difference in utility when there is no loss, $\{(1 - p_b)[U(Y - \pi) - U(Y)]\}$, and the difference in utility when there is a loss, $\{p_b E_m[U(Y - \pi - \theta(m)) - U(Y - m)]\}$. Notice that, if there is no loss, utility is lower if insurance has been purchased than if it has not been purchased since a premium was paid with no recovery. If there is a loss, however, expected utility will be greater with insurance than without. The insurer with $p_b = p_b^*$ is the firm where these two terms exactly balance each other.

Now consider any $p_b > p_b^*$. For each of these firms, the weight placed on expected utility when there is a loss is greater, and the weight placed on utility without a loss is smaller, than for the firm with $p_b = p_b^*$. Thus, each of these firms will find reinsurance more attractive than does the firm with $p_b = p_b^*$. Therefore, all firms with $p_b > p_b^*$ will purchase reinsurance, and all firms with $p_b < p_b^*$ will not purchase reinsurance.

Combining equations (A1), p_b^* is determined as

$$(A2) \quad p_b^* = \frac{U(Y) - U(Y - \pi)}{U(Y) - U(Y - \pi) + E_m[U(Y - \pi - \theta(m)) - U(Y - m)]}.$$

The reinsurance premium π^* is the average reinsurance payment conditional on having a loss (the loss less the primary insurer's payment) times the probability that a firm purchasing insurance has a loss. This is given by

$$(A3) \quad \pi^* = \{E_m[m] - E_m[\theta(m)]\} \cdot E[p_b | p_b > p_b^*].$$

The optimal reinsurance policy maximizes expected utility over p_b . If there were only one p_b , full insurance would be optimal ($\theta[m] = 0$). With differing values of p_b , however, only the firms with higher p_b will purchase insurance.

To combat adverse selection, think about increasing $\theta(m)$ on the first dollar of losses. As this is done, the premium falls by the expected primary-insurer payment. Starting from the point of full insurance, this reduction in coverage has no welfare cost for firms already purchasing reinsurance. For firms at the margin of purchasing reinsurance, however, the reduction in reinsurance coverage is attractive; reinsurance is actuarially overpriced for this group. Thus, some additional firms will choose to purchase reinsurance. This reduces the

reinsurance premium by even more, benefiting the firms that were already purchasing insurance. Thus, incomplete reinsurance will be optimal.

To design the optimal reinsurance policy, consider a given amount of total cost sharing in the reinsurance policy, $E_m[\theta(m)]$. The direct effect of this cost sharing on the reinsurance premium (eq. [A3]) is independent of where in the distribution of m this cost sharing occurs. But selection and welfare of insurers purchasing reinsurance depends on how this cost is determined. Note that firms purchasing reinsurance want to design cost sharing to maximize $E_m[U(Y - I - \theta(m))]$. Maximizing this value also reduces p_b^* the most (eq. [A2]).

Thus, the goal is to design the reinsurance policy that maximizes $E_m[U(Y - I - \theta(m))]$ over all the policies with the same level of total cost sharing $E_m[\theta(m)]$. Given risk aversion, the most valuable insurance to provide is that for the greatest risks; the next most valuable insurance is that for the next highest risk; and so forth. Thus, the optimal policy must insure all the high risks up to the point where the budget constraint is met. This is exactly the excess-of-loss policy described in the text.

References

- Akerlof, George. 1970. The market for lemons. *Quarterly Journal of Economics* 84: 488–500.
- Angell, Marcia. 1996. *Science on trial: The clash of medical evidence and the law in the breast implant case*. New York: Norton.
- A. M. Best Co. 1996. *Best's aggregates and averages—property-casualty*. New York.
- Borden, Sara, and Asani Sarkar. 1996. Securitizing property catastrophe risk. *Current Issues in Economics and Finance* (Federal Reserve Bank of New York) 2 (August): 1–6.
- Cochrane, John. 1995. Time consistent health insurance. *Journal of Political Economy* 103:445–73.
- Conning & Co. 1993. *Reinsurance: Too little or too much?* New York.
- Cutler, David M. 1996. Why don't markets insure long-term risk? Harvard University, Economics Department. Mimeo.
- Giles, Martin. 1994. A survey of insurance. *Economist*, 3 December, suppl., 3–22.
- Huber, Peter. 1988. *Liability: The legal revolution and its consequences*. New York: Basic.
- Insurance Information Institute. 1996a. *Catastrophes: Insurance issues*. New York.
- . 1996b. *Environmental pollution: Insurance issues*. New York.
- . 1996c. *Reinsurance*. New York.
- Kaplow, Louis. 1991. Incentives and government relief for risk. *Journal of Risk and Uncertainty* 4:167–75.
- Litzenberger, Robert H., David R. Beaglehole, and Craig E. Reynolds. 1996. *Assessing catastrophe-reinsurance-linked securities as a new asset class*. New York: Goldman Sachs.
- Major, John. 1996. Basis risk and catastrophe reinsurance. Guy Carpenter & Co. Mimeo.

- Myerson, Roger. 1979. Incentive compatibility and the bargaining problem. *Econometrica* 47:61–73.
- Pratt, John W., and Richard J. Zeckhauser. 1989. The impact of risk sharing in efficient decision. *Journal of Risk and Uncertainty* 2:219–34.
- Rothschild, Michael, and Joseph Stiglitz. 1976. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Quarterly Journal of Economics* 90:629–50.
- Viscusi, W. Kip, Richard J. Zeckhauser, Patricia Born, and Glenn Blackmon. 1993. The effect of 1980s tort reform legislation on general liability and medical malpractice insurance. *Journal of Risk and Uncertainty* 6:165–86.
- Weiler, Paul C., Howard H. Hiatt, Joseph P. Newhouse, William G. Johnson, Troyen A. Brennan, and Lucian L. Leape. 1993. *A measure of malpractice: Medical injury, malpractice litigation, and patient compensation*. Cambridge, Mass.: Harvard University Press.
- Wilson, Robert B. 1968. On the theory of syndicates. *Econometrica* 36:119–32.
- Zeckhauser, Richard J. 1991. The strategy of choice. In *Strategy and choice*, ed. Richard J. Zeckhauser. Cambridge, Mass.: MIT Press.

Comment John H. Cochrane

This paper is an ambitious survey. It is full of interesting anecdotes, facts about reinsurance, and speculations about forces that might move reinsurance contracts in one direction or another. While much time could pleasantly be spent going through the paper and discussing each issue that it raises, instead I will concentrate on the central analytic point. Here's the setup: m denotes the size of the loss if there is one, with probability $g(m)$. The probability that any loss occurs is p . This probability is random, too, and $f(p)$ is the “probability of probability” p . The original insurer may have better information about the probability of loss $f(p)$ or about the distribution of the size of losses $g(m)$.

I emphasize that the issue here is, not the probability structure itself, but *differences* between the insurer's knowledge and the reinsurer's knowledge. The fact that we are all uncertain about what probabilities to attach to cataclysmic events has nothing to do with adverse selection.

The central point is stated right up front: “When the asymmetry is on the probability of a loss [$f(p)$] but not its magnitude [$g(m)$], the optimal reinsurance contract is a standard ‘excess-of-loss’ policy [deductible]. . . . When the asymmetry is over the magnitude of the loss [$g(m)$], the optimal reinsurance policy covers smaller losses as well as large losses but leaves the primary insurer exposed for some large losses [risk sharing].” I think that this proposition is a misleading summary of what we learn from the analysis in the paper.

John H. Cochrane is the Sigmund E. Edelstone Professor of Finance at the Graduate School of Business, University of Chicago; a consultant at the Federal Reserve Bank of Chicago; and a research associate of the National Bureau of Economic Research.

All that can matter to the optimal contract is each player's probability distribution over final outcomes—losses m in this case. Since there are no intermediate moves, the sequence of signals must be irrelevant. Thus, suppose that you get to distributions over final outcomes by different values for p . You obtain *exactly the same optimal contract* if you get to those distributions over final outcomes by different values for $g(m)$. The labels *uncertainty over probabilities* and *uncertainty over losses* must be meaningless in terms of the optimal contracts that they generate.

What about Cutler and Zeckhauser's examples?

Both the "asymmetric information about probability" example and the "asymmetric information about losses" example generate situations in which the reinsurer knows that he faces an original insurer that has one of two probability distributions over final events. He has to design a clever schedule of payments as a function of losses to separate out original insurers with one versus the other type of probability distribution.

The examples differ in what the two types of final distributions over outcomes are. It should not be a surprise that contracts designed to separate people with one pair of probability distributions over final outcomes are different from contracts that separate people with much different pairs of such distributions. In fact, the paper has a very nice exposition of the intuition for this result: "If the buyer knows more about the probability of a loss than its magnitude, reinsurance should exclude coverage for the common small risks that are the source of exploitable adverse selection; this is exactly the 'deductible-style' policy that characterizes excess-of-loss insurance. If the buyer knows more about the magnitude of loss than its probability, however, it is optimal to make it costly for him to buy reinsurance when he knows that the magnitude of loss is high."

This intuition is right on the money, but notice that it concerns only the uncertainty that reinsurers have about the original insurers' probabilities over *final outcomes*, not how one gets there. Another way to put the point is that, in general, the uncertainty over probabilities and the uncertainty over distributions are not separately identified by the uncertainty over distributions over final outcomes. One can tell the difference between these two only in highly structured and stylized examples. In this case, the crucial "identifying assumption" is that $g(m)$ places no weight on $m = 0$. Absent that identifying assumption, we can rewrite the uncertainty in probability example as an uncertainty in distribution example, and you would never know the difference between the two.

These are also perfectly standard results in the theory of insurance—for example, Rothschild and Stiglitz (1976) or Kreps (1990, 668). These days, one usually finds optimal contracts with the "revelation principle" rather than by the supply-and-demand methods used in this paper (or at least the conference draft). A risk-neutral insurer offers a contract to a risk-averse consumer. The consumer is one of two or more "types" with different probability distribution

of losses. To find the optimal contract—the optimal payment given each possible loss—we maximize the insurer’s profits subject to the constraints that each type of consumer will accept the contract he is offered and that each type of consumer will correctly reveal his type. In equations, one solves problems like the following:

PROBLEM. Choose state-contingent payment to maximize insurer profit

$$\max_{\{q(m_i, \cdot)\}} \Pr(A) \sum_i \pi_i^A q(m_i, A) + \Pr(B) \sum_i \pi_i^B q(m_i, B)$$

(where $\Pr[A]$ = insurer’s probability that consumer is type A, m_i = loss in state i , π_i^A = type A probability of loss m_i , and $q[m_i, A]$ = net payment if consumer says type A and loss is m_i) subject to (1) type A consumers accept the contract:

$$\sum_i \pi_i^A u[w - m_i - q(m_i, A)] \geq \sum_i \pi_i^A u(w - m_i);$$

(2) type B consumers accept the contract:

$$\sum_i \pi_i^B u[w - m_i - q(m_i, B)] \geq \sum_i \pi_i^B u(w - m_i);$$

(3) type A consumers do not pretend that they are type B:

$$\sum_i \pi_i^A u[w - m_i - q(m_i, A)] \geq \sum_i \pi_i^A u[w - m_i - q(m_i, B)];$$

(4) type B consumers do not pretend that they are type A:

$$\sum_i \pi_i^B u[w - m_i - q(m_i, B)] \geq \sum_i \pi_i^B u[w - m_i - q(m_i, A)].$$

We know how to solve problems like this and how to find *optimal* contracts. The optimal contracts behave just as Cutler and Zeckhauser find in their examples. There will be less insurance—deductibles, caps, copayments, etc.—in the range of outcomes where the possible probabilities differ the most or “odds ratios” are highest. In fact, one can reverse engineer problems like this to find an assumption about the nature of private information to generate any desired sharing rule or contract.

In sum, these examples are perfectly standard insurance theory. They have nothing special to do with reinsurance or cataclysms, and the results—the degree to which different sizes of loss are covered—depend only on probability distributions over final outcomes, not in any fundamental way whether those distributions are generated by “uncertainty over probabilities” or “uncertainty over losses given there is a loss.” If one wants to avoid revelation-principle constructions, then either one has to find a new, restricted definition of optimal contract, or one is just showing a different implementation mechanism for a revelation-principle contract.

I do not mean to sound critical. I think that it is a good thing that the analysis in this paper comes down to perfectly standard theory. Thousands of papers have been written on optimal contracting with asymmetric information. It

would be astounding if this paper came up with an original theory. I much *prefer* papers that apply standard theory to interesting practical situations. It gives me comfort as an economist that we do not have to invent a new and clever theory for every situation. I am in the same boat. The paper I wrote (Cochrane 1995) that got me invited to this conference applied contract theory so old-fashioned there was not even private information. The point of the above forest of equations is constructive. It is a suggestion for how Cutler and Zeckhauser can actually compute optimal contracts in the situation that they analyze, rather than analyzing the relative merits of arbitrary contracts.

But, while there have been thousands of papers working out the theory, there have been almost no papers that try to see whether this beautiful theory has anything to do with our world. Is there really any asymmetric information, and what form does it in fact take? *Let us* reverse engineer-observed real-world contracts and see if the required uncertainty about insurer's probability distributions over losses makes any sense. *Let us* examine what information we have on actual insurer's probability distributions. Along this path, completing the marriage of theory and detailed industry knowledge toward which I see this paper working, I think that Cutler and Zeckhauser can end up with something really useful.

This said, I must voice my skepticism. Of the standard textbook problems with insurance contracts, of all the impediments to cataclysm reinsurance, is adverse selection—symmetric information over the probability distribution of losses—really the most important?

Adverse selection is already not so obvious for people. Do you really know more about your health on the basis of your aches and pains than a doctor armed with your medical history and actuarial tables? For institutions—insurance companies selling to reinsurers—symmetric information seems even more tenuous. Anything that an institution knows and can act on is written down somewhere and can be shared with a reinsurer.

Furthermore, the uncertainties about the probability of cataclysmic losses that many of us have commented on during this conference argue to me against asymmetric information as a serious problem. The heart of asymmetric information problems is that the reinsurer does not want to offer a contract and have only those original insurers who *know* the probabilities are a lot worse than it thinks they are take the contract. It is not enough that the insurers happen to *believe* that the probabilities are a lot worse: so long as the insurers do not *know* a lot more than the reinsurer does, the reinsurer will not discover ex post a surprisingly large exposure, and contracts will work fine. If nobody knows the right probabilities, there will be little information and lots of confusion as to correct pricing. But there will not be *asymmetric* information and adverse selection.

We can think of testing for adverse selection rather than intuiting or investigating the extent of asymmetric information. I just read an interesting paper (Cawley and Phillipson 1996) that tested whether asymmetric information is

important in life insurance. The theory, like that presented here, predicts that it should be cheaper to buy a little insurance than a lot of insurance, to keep really sick people from loading up on insurance. Insurers should make sure that consumers do not undo this size premium by getting insurance from several different companies. But, in fact, there are volume discounts, and insurers do not care if you take out multiple policies! Cawley and Phillipson also find that insurers, with good actuaries and models, have *better* information than customers on likely life spans!

So bring on the empirical work. Cutler and Zeckhauser seem ideally ready to do a similar exploration of the catastrophe-reinsurance market. But, as with life insurance, let us be prepared to find that, although adverse selection fascinated a generation of theorists, it might not have much to say about many actual insurance contracts.

A word of caution, however. Adverse selection leads to nonlinear risk sharing, but not every nonlinear contract means adverse selection. For example, consider a company that is hedging oil-price or interest rate risk in derivatives markets. It will sign many contracts with highly nonlinear payouts, payouts that look like “deductibles,” “caps,” “slices,” and so forth. But, obviously, the structure of these securities has *nothing* to do with adverse selection. Analogously, I think that it could be very dangerous to jump to an adverse-selection interpretation of nonlinear contracts in an increasingly securitized reinsurance market.

Of the textbook asymmetric information stories, I would guess that asymmetric information about how big the losses really are *once they have happened* is a far bigger problem. You do not just call up your reinsurer and say, “It rained a bit. Send us that check for \$100 billion.” The contracts must be carefully written and triggered on events that can be easily verified and not subject to manipulation. Currently, reinsurance and even the exchange-traded cat options are triggered on reports of industry losses. To a cynic like me, this places an extraordinary amount of faith in accountants, adjusters, and reporting services. But, by focusing attention on this issue and giving an incentive for the production of good indices, the process of securitization seems likely to help.

I would also follow much of the discussion here and guess that loss reserves, liquidity, capital, etc. are crucial issues. David Cutler asked the question, “Who should pay?” The answer is that we all should pay. The (Pareto-)optimal response to a large shock is that everyone should pitch in a bit. The question is, What institutions or securities best implement this result?

We worry that cataclysms, although trivial compared to national wealth, nonetheless can exhaust reserves, setting off a chain of bankruptcies. Yet tying up more reserves in liquid assets—building up capital to cover the largest cataclysm—is obviously inefficient.

The central problem is limited liability. The answer, it seems to me, is more widespread use of securities that allow some liability. Each \$10,000 of a pension fund should include a security, or just a contractual arrangement, that costs

nothing, gets, say, \$5.00 per month premium, but exposes the fund to a liability of up to, say, \$100 in cataclysm or calamity events. The current pricing of catastrophe reinsurance means that these securities earn very large expected returns even though they are uncorrelated with other security returns. One could come close to this arrangement with the cat options that are already traded.

This arrangement allows us all to share in the risk, which is still, after all, a small fraction of total wealth. It does so without the need for huge inefficient loss reserves or the distortions and inefficiencies of implicit government catastrophe insurance by *ex post* taxation. The design, implementation, and marketing of such securities seems to me the most promising response to the issue of cataclysm reinsurance. Of course, pension funds must be allowed to hold such securities, so, in the end, regulation may be the reason that we do not already all share in catastrophe risk, and regulatory change may be an important step in allowing us to do so.

References

- Cawley, John, and Tomas Phillipson. 1996. An empirical examination of information barriers to trade in insurance. Working Paper no. 5669. Cambridge, Mass.: National Bureau of Economic Research.
- Cochrane, John H. 1995. Time-consistent health insurance. *Journal of Political Economy* 103:445–73.
- Kreps, David M. 1990. *A course in microeconomic theory*. Princeton, N.J.: Princeton University Press.
- Rothschild, M., and J. Stiglitz. 1976. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Quarterly Journal of Economics* 90:630–49.