

The Pricing of U.S. Catastrophe Reinsurance

Kenneth A. Froot and Paul G. J. O'Connell

The price of catastrophe reinsurance in the United States has fluctuated markedly in recent years. These fluctuations are commonly associated with the pattern of catastrophe occurrences. For example, catastrophe losses during the period 1992–94 totaled \$38.6 billion in 1994 dollars, exceeding the cumulative total of losses during 1949–91 of \$34.6 billion. During this three-year period, prices on catastrophe-reinsurance cover more than doubled and then began to decline thereafter. What drives such changes in price? Does the demand for reinsurance shift, does the supply of reinsurance capital change, or do both occur?¹

If catastrophe losses lead to a decrease (leftward shift) in supply, then we would expect to see increases in price coupled with declines in quantity after an event. Of course, a decline in supply is possible only in the presence of some form of capital market imperfection. If capital markets were perfect, the supply curve for reinsurance would be perfectly elastic. In this case, regardless of losses, the price of reinsurance would be fixed, where the “price” of a con-

Kenneth A. Froot is the Industrial Bank of Japan Professor of Finance and director of research at the Harvard Business School and a research associate of the National Bureau of Economic Research. Paul G. J. O'Connell is president of Emerging Markets Finance, LLC.

The authors are grateful to Peter Diamond, Chris Milton, Julio Rotemberg, Jeremy Stein, Rene Stulz, and conference participants for helpful comments; to David Govrin, Chris McGhee, Denis Medvedsek, Brian Murphy, and Barney Schauble for help in obtaining and understanding reinsurance data; and to the NBER and the Division of Research at Harvard Business School for generous research support. All errors are the authors'.

1. A number of papers have investigated these cycles, attempting to identify supply and demand shocks in insurance markets. Cummins and Outreville (1987) show that lags in data collection or price regulation can generate cycles in property-casualty underwriting margins. Gron (1994) presents evidence that, assuming that there are no marketwide demand shocks, the cycles in property-casualty margins are due to variation in the supply of insurance capacity rather than institutional lags or reporting practices. Gron and Lucas (1995) investigate why these cycles appear to be so persistent. They find that, when the net worth of insurers declines, the total amount of capital raised through security issues is small. See also Winter (1988) and Cummins and Danzon (1991).

tract is best thought of as the ratio of premiums to actuarially expected losses covered under that contract. Capital market imperfections would imply that the marginal cost of producing reinsurance is increasing in the quantity supplied. Thus, these imperfections lead to an upward-sloping supply curve, which (all else equal) can shift back as a result of reinsurer losses. Such a supply shift increases price and reduces quantity. As with price, it is best to think of this "quantity" as the actuarially expected loss covered by reinsurance.

On the other hand, catastrophe losses may lead to increases in demand. Rightward demand shifts can be thought of as the result of an actual or perceived increase in actuarial losses covered by a given contract. We call this *probability updating*. Naturally, it would seem possible to identify such demand shifts from the fact that they lead to an increase in price *and* quantity. Thus, conditional on a loss, an absolute decline in the quantity of reinsurance purchased would be evidence of important leftward shifts in supply even if there were also positive increases in demand.

We look for such absolute declines in quantity, but, in addition, we pursue the probability-updating hypothesis further. While it is impossible to distinguish between probability updating and capital market imperfections on the basis of the behavior of aggregate price indices over time, it is possible to distinguish between them on the basis of the behavior of *cross-sectional* changes in reinsurance prices. Specifically, probability updating ought to vary across contracts, with larger price increases associated with contracts for which more probability updating occurs. We therefore examine cross-sectional price increases in response to an event and determine the extent to which they are explained by relative contract exposures.

To see how this works, consider a catastrophe loss caused by a winter freeze in New England. We might expect such a loss to affect strongly (and positively) the distribution of prospective losses due to freeze and/or the distribution of prospective losses due to other perils in New England. After all, the event may cause people to recognize how much damage a freeze can do or to learn about the replacement costs of certain physical assets in New England. However, such updates in knowledge would have little or no import for the distribution of catastrophe losses outside New England, where freezes do not occur. Specifically, little would be learned about loss exposures in California (which faces primarily earthquake risk), the Southeast (which faces primarily hurricane risk), or Texas (which faces primarily windstorm risk). Under probability updating, it follows that contracts with relatively little exposure to freeze and/or to the Northeast region ought to have relatively small price increases. In this way, we are able to further distinguish between capital market imperfections and probability updating.

Our identification strategy is made possible through the use of a unique and detailed data set from Guy Carpenter and Company, by far the largest catastrophe-reinsurance broker for U.S. catastrophe exposures. These data include all U.S. catastrophe-reinsurance contracts brokered by Guy Carpenter between

1970 and 1994. They allow us to measure prices and contract losses and to go about the complex process of estimating each contract's exposure to different event types and regions.

To preview our results, we find that supply, rather than demand, shocks are more important for understanding the effect of losses on reinsurance prices and quantities. Capital market imperfections therefore appear to be the dominant explanation. There is limited evidence for probability updating, and what evidence there is suggests that the effect is of a small magnitude. The magnitudes of the supply effects are large: after controlling for relative contract exposure, a \$10 billion catastrophe loss raises average contract prices by between 19 and 40 percent and reduces quantity of reinsurance purchased by between 5 and 16 percent.

The rest of the paper is organized as follows. Section 5.1 sets out our identification strategy and the structure of our empirical tests. Section 5.2 describes our data sources. In section 5.3, the calculation of contract exposure and price is discussed in detail. We devote considerable attention to the calculation of exposure, which requires a number of involved steps. Section 5.4 provides a brief graphic analysis. The empirical testing is carried out in section 5.5. Section 5.6 summarizes and offers our conclusions.

5.1 The Price and Quantity of Reinsurance

We examine the equilibrium prices and quantities of single-event excess-of-loss reinsurance contracts. These contracts help reinsure insurance companies against losses resulting from natural catastrophes in the United States, such as windstorms or earthquakes.

To understand how such contracts work, consider an insurer that purchases a layer of reinsurance covering \$100 million of losses in excess of \$200 million. These terms imply that, if the insurer's losses from a single catastrophic event during the contract year exceed \$200 million (the "retention"), the layer is triggered. The reinsurer pays the insurer the amount of any losses in excess of \$200 million, with the maximum payment—the "limit"—capped at \$100 million.² By purchasing this contract, the insurer cedes its exposure to single-event catastrophe losses in the \$200–\$300 million range. In return for assuming this exposure, the reinsurer receives a payment, known as the "premium." If the insurer wishes to cede a broader band of exposure, it could purchase additional layers—\$100 million in excess of \$300 million, \$100 million in excess of \$400 million, and so on.

The price of a reinsurance contract is best measured as the premium per unit of exposure. In the marketplace, premium is usually expressed relative to limit (the ratio is called *rate on line*). However, limit is a poor proxy for contract

2. To guard against moral hazard, excess-of-loss reinsurance contracts typically require coinsurance. In practice, this effectively means that the insurer provides 5–10 percent of the reinsurance.

exposure—it ignores the level of the contract's retention, for example. To remedy this, we measure price by premium per unit of actuarially expected loss. Indeed, we use actuarially expected loss as our measure of the “quantity” of reinsurance purchased.

In section 5.3, we describe how historical data on catastrophe losses and company-specific market-share information can be used to measure the actual exposure of each reinsurance contract. It is important to emphasize that we measure actuarially expected loss from a loss distribution that is *time invariant*. We cannot condition our measures of expected loss on previous losses. To the extent that loss distributions shift in response to recent loss history, we measure both quantity and price with error. These potential mismeasurements are important for the way we specify our tests and hypotheses.

5.1.1 Identifying Capital Market Imperfections

As noted above, we investigate two channels by which catastrophe losses affect reinsurance prices and quantities. First, “capital market imperfections” may impede the flow of capital into the reinsurance sector. There may be several sources of such imperfections. One is that existing reinsurers may find it costly or undesirable to *raise* additional external capital. These costs could result from information asymmetries between managers and owners (which implies that equity-sale announcements drive down share prices) or from dilution of managerial control (which implies that managers are averse to expanding the capital base). Another potential source of imperfection is that it is costly to *carry* equity capital. These costs may accrue from forgone tax shields, agency problems, or, in the case of reinsurers, frictional collateral costs.³

If we could accurately measure, at each point in time, the distribution of one-year-ahead losses conditional on all information, then it would be relatively simple to test the capital market imperfections story. In the presence of such imperfections, capital depletion associated with event losses constricts the supply of reinsurance, driving up the price of all contracts. If, for example, the event were a hurricane, the supply of hurricane-reinsurance capacity would fall. So would the supply of nonhurricane-reinsurance capacity since both exposures are borne by the same capital base. Figure 5.1 shows the effect on equilibrium prices and quantities for contracts that are exposed to hurricane risk (fig. 5.1A) and those that are not (fig. 5.1B). Note that the exposure supply curves are upward sloping. This is due to the capital market imperfections, which raise the marginal cost at which reinsurers are able to offer successively greater exposure protection to insurers.⁴

Figure 5.1 suggests that capital market imperfections generate a negative correlation between prices and quantities since loss shocks lead to shifts in the

3. For a survey of these costs and their effect on financing patterns, see Froot (1995). Froot and Stein (1998) study the implications for financial intermediaries of costly equity finance.

4. For a model of the reservation price that a financial intermediary such as a reinsurer is willing to offer marginal units of risk exposure, see Froot and Stein (1998).




Fig. 5.1 Capital market imperfections: A, Hurricane-exposed contracts; B, Non-hurricane-exposed contracts




Fig. 5.2 Real quantity of catastrophe exposure ceded, 1975:1–1993:4

Note: In any quarter, the series is the sum of the contract exposure that is ceded by four major insurers in that quarter. The four companies represent approximately 10 percent of the market.

supply curve. If losses lead to demand-curve shifts, on the other hand, a positive correlation between price and quantity results.

Since negative correlation would seem to be evidence in favor of supply shifts (and, therefore, capital market imperfections), it seems useful to ask whether prices and quantities are in fact negatively correlated. Figures 5.2 and 5.3 show our measures of quantity and price, respectively, for the sample period for which we have contract data.⁵ More specifically, figure 5.2 shows an

5. For a description of our data and the computation of prices and quantities, see sec. 5.2 below.




Fig. 5.3 Industry price per unit of ceded exposure, 1975:1–1993:4

index of the total quantity of catastrophe exposure that was ceded in the U.S. reinsurance market from 1975:1 to 1993:4. The series is calculated by summing all the exposure embodied in the excess-of-loss contracts in each quarter and dividing it by the total market share represented by the contracts. Figure 5.3 plots the industry price series quarterly from 1975:1 to 1993:4. Each observation is the exposure-weighted average of the price of all contracts that are in force in that quarter.⁶

A number of features of these figures are noteworthy. First, it appears that quantities rose and prices fell for much of the late 1970s and 1980s. Second, a startling rise in prices and decline in quantities took place beginning in the mid-1980s through the end of the sample period. Indeed, in 1993, price was between five and seven times its historical average. This will come as no surprise to industry observers. It is common to relate this price rise to the occurrence of a number of large events during this period, notably Hurricane Andrew (\$20 billion in losses) in August 1992, Hurricane Hugo in 1989, and several windstorms in 1985–86. Figure 5.4, which plots total catastrophe losses by quarter from 1970:1 to 1994:4 as measured by Property Claims Ser-

6. These industry series are based on the contract prices and exposures for four insurers that purchased reinsurance through Guy Carpenter in every year from 1975 to 1993. The series are representative of the behavior of prices and quantities for the other insurers in our database.




Fig. 5.4 Total adjusted PCS losses by quarter, 1970:1–1994:4

vices, lends support to this view.⁷ In the period since 1994 (for which we have no data), the price of reinsurance has declined and quantity increased somewhat, notwithstanding the occurrence of the Northridge Earthquake in January 1994. From these observations, it is clear that there is considerable negative correlation between prices and quantities at frequencies of several years.

Figures 5.5 and 5.6 provide further evidence of the apparent negative correlation between equilibrium prices and quantities. Figure 5.5 plots industry price-quantity pairs. It is interesting because it suggests the existence of two regimes in catastrophe-reinsurance pricing. The 1970s and early 1980s saw strong expansion in the quantity of risk ceded coupled with a moderate decline in per unit prices. The late 1980s and early 1990s were characterized by ballooning prices coupled with quantity declines. One interpretation of these patterns is that the 1970s and 1980s were a period of expanding reinsurance demand while the 1990s exhibit a contraction of reinsurance supply. Figure 5.6 plots each contract's price against its exposure. Both variables have been demeaned by insurer. As indicated by the linear fit to the plot (slope -0.32 , standard error 0.04), when a contract embodies less-than-average exposure, it tends to be priced above average. Taken at face value, the points appear to lie on firm-specific reinsurance-demand schedules. Or, put another way, strong

7. For a description of this loss series, see sec. 5.2 below.




Fig. 5.5 Industry price-quantity pairs, 1975:1–1993:4




Fig. 5.6 Contract-by-contract price-quantity pairs, demeaned by company, 1970–94 (linear fit estimated by OLS)




Fig. 5.7 Probability updating: *A*, Hurricane-exposed contracts; *B*, Non-hurricane-exposed contracts

evidence of negative correlation between price and quantity suggests that shifts in the reinsurance supply curve have been important.

Taken by itself, this evidence would seem to support the hypothesis of capital market imperfections. However, it is important to see the implications of demand shocks, or what we call *probability updating*. Probability updating holds that the occurrence of a catastrophe may raise the real or perceived distribution of losses above what we measure with our time-invariant loss distribution. For example, after Hurricane Andrew in 1992, some insurers were surprised that the construction methods used for houses in Homestead, Florida, performed so poorly in high winds. Andrew might also have led to upward shifts in agents' subjective distributions.

To see the effects of probability updating, suppose that, after a particular event, say, a hurricane, agents update positively about the likelihood of hurricane losses. Then, even if the premium per unit of actual exposure stays constant, the ratio of premium to *observed* exposure rises. This is because we measure exposure from a time-invariant loss distribution. Figure 5.7 traces this effect on supply and demand. Suppose that capital markets are perfect—the supply of capital to the reinsurance sector is infinitely elastic at a given price. In the aftermath of the hurricane, the observed supply curve shifts upward, as shown in figure 5.7A. To understand what happens to the demand schedule, consider an insurer that, at a given price, wants to cede the same amount of exposure before and after the hurricane. Since the perceived risk of hurricanes increases, the terms of this insurer's contract must be rewritten to keep exposure constant. For example, the retention on the contract could be raised. Any such change that keeps *actual* ceded exposure constant causes *measured* exposure to fall and, concomitantly, *measured* price to rise. Thus, the entire observed demand schedule shifts upward and to the left. All this assumes that there is no actual change in demand. However, the demand curve may exhibit further changes as agents alter the amount of reinsurance they are willing to demand at a *given* price. For example, the demand schedule may shift out if




Fig. 5.8 Capital market imperfections and probability updating: *A*, Hurricane-exposed contracts; *B*, Non-hurricane-exposed contracts

homeowners learn from the hurricane that they are underinsured. Alternatively, the schedule may shift in if coastal homeowners who sustain severe damage from the hurricane elect to move to a less-exposed inland region rather than to rebuild in the same place.

Consequently, the overall effect on the observed demand curve is ambiguous. If the observed DD curve experiences no net shift at all, the equilibrium moves from point A to point C , which results in an increase in price and a decrease in demand. Thus, demand shifts combined with probability updating can duplicate the finding that prices increase while quantities decrease. We therefore cannot conclude that such negative correlation is evidence of capital market imperfections. On the other hand, if the observed DD curve shifts out on net, as shown in the figure, the equilibrium shifts from point A to point B .

Probability updating, then, complicates the process of identifying capital market effects. To overcome this problem, consider the following identification strategy. Assume that hurricane losses may lead to probability updating for hurricanes but not for other types of events in other regions. The idea here is simple: if a hurricane occurs, it may not alter the perceived frequency or severity of nonhurricane events. Accordingly, probability updating from a hurricane should produce no change in prices or quantities of contracts that have little or no hurricane exposure. However, if there are capital market imperfections, then prices tend to increase and quantities decrease on nonhurricane-exposed contracts.

Figure 5.8 illustrates. A hurricane loss that raises the perceived likelihood of future hurricane losses will simultaneously shift the observed demand and supply curves upward and to the left. These effects are shown as movements in the schedules from DD to DD' and from SS to SS' in figure 5.8A, which together produce an increase in price. There are no such probability-updating effects for the contracts that are not exposed to hurricane risk in figure 5.8B. However, if the catastrophe loss depletes the pool of capital that is available to the industry for all types of reinsurance, then the supply schedule for both

types of contracts shifts back to SS'' . In these circumstances, we can expect the price of both to increase.

A natural way to implement this identification scheme is to regress changes in price and quantity on lagged losses by type and lagged losses by type interacted with contract exposure by type. If there are K types of catastrophe, the conditional expectation functions takes the form

$$(1) \quad \Delta \ln(p_{j,t}) = \alpha + \sum_{k=1}^K \beta_k \theta_t(l_k) + \sum_{k=1}^K \gamma_k w_{j,k,t} \theta_t(l_k) + \varepsilon_{j,t},$$

$$(2) \quad \Delta \ln(q_{j,t}) = \alpha + \sum_{k=1}^K \delta_k \theta_t(l_k) + \sum_{k=1}^K \phi_k w_{j,k,t} \theta_t(l_k) + \varepsilon_{j,t},$$

where $p_{j,t}$ and $q_{j,t}$ are the price and quantity, respectively, of insurer j 's contract at time t , $\theta_t(l_k)$ is a distributed lag of losses of event type k , $w_{j,k,t}$ is the exposure of insurer j to catastrophes of type k at time t , and $\varepsilon_{j,t}$ is a disturbance term. The relative exposure term, $w_{j,k,t}$, is defined as

$$(2) \quad \Delta \ln(q_{j,t}) = \alpha + \sum_{k=1}^K \delta_k \theta_t(l_k) + \sum_{k=1}^K \phi_k w_{j,k,t} \theta_t(l_k) + \varepsilon_{j,t},$$

where $q_{j,k,t}$ is the absolute exposure (the actuarially expected loss) of insurer j to catastrophes of type k at time t . β_k (δ'_k) measures the response of contract prices (quantities) to losses of type k independent of contract exposure. The parameter γ_k (ϕ'_k) detects the exposure sensitivity of the price (quantity) responses. In estimating (1) and (2), we restrict attention to simple lag structures, such as⁸

$$(4) \quad \theta_t(l_k) = l_{k,t-2} + l_{k,t-3} + l_{k,t-4} + l_{k,t-5}.$$

To the extent that capital market imperfections are responsible for price and quantity changes, we expect $\beta_k > 0$ in (1) and $\delta_k < 0$ in (2), with the null hypothesis of perfect capital markets being $\beta_k = 0$ and $\delta_k = 0$. Conditional on a loss from a particular event type, prices tend to rise and quantities to fall equally across all contracts if there are capital market imperfections. The composition of contract exposure matters for price changes only if probability updating is present. Thus, with no probability updating, $\gamma_k = 0$ for all k , and, with probability updating, $\gamma_k > 0$ for all k . Of course, if both capital market imperfections and probability updating are important, we expect $\beta_k > 0$ and $\gamma_k > 0$.

It is of interest to consider several modifications of the basic specifications (1) and (2). For instance, if the average effect of a dollar of losses is the same across catastrophe types, as the capital market imperfections view would suggest, then we expect $\beta_1 = \beta_2 = \dots = \beta_K$. This restriction can readily be im-

8. Losses in the quarter immediately preceding contract inception appear to have no influence on prices. The most likely reason for this is that there can be a delay in assessing the extent of catastrophe losses. Accordingly, losses lag prices by two or more quarters in all our analyses.

posed in the estimation. A second worthwhile modification is to allow the average and marginal effects of losses to differ. If the supply schedule is nonlinear, for example, then it may be true that large losses have proportionately bigger effects than small losses. Indeed, this is true in most models of capital market imperfections, as discussed in Froot and O'Connell (1997). A convenient way to allow for this possibility is to include higher-order terms in the distributed lag of losses.

5.1.2 Differentiating among Catastrophe Types

We differentiate exposure along two dimensions: by geographic region and by catastrophe type. Distinguishing events across regions is appealing because regional market-share data are available for many reinsurance contracts. These market-share data provide a ready measure of relative contract exposure. Thus, $w_{j,k,t}$ can be defined as

$$(5) \quad w_{j,k,t} = \frac{m_{j,k,t}}{\left(\frac{1}{K}\right) \sum_{k=1}^K m_{j,k,t}},$$

where $m_{j,k,t}$ is the market share of insurer j in region k . We estimate (1) and (2) using this measure of exposure in the empirical analysis in section 5.5 below.

There are two drawbacks to this approach, however. First, market shares are typically built up from the share of total catastrophe premiums in a region that go to insurer j . To the extent that premiums are a poor proxy for true exposure, $w_{j,k,t}$ will be subject to measurement error. Second, it is clear that the catastrophe occurrences in one region may reveal important information about the distribution of losses in other regions. For example, hurricane losses in the Southeast are likely to lead to probability updates in the distribution of hurricane losses in both the Northeast and the Southeast. As a result, price increases (and quantity decreases) in the Northeast in the aftermath of a Florida hurricane could be due to either probability updating or inelastic capital supply to the industry, and identification may be tenuous.

These two shortcomings motivate us to differentiate catastrophes by type as well. Here, the exposure of each contract to distinct classes of events, such as earthquakes, hurricanes, and winter storms, is calculated. This strategy avoids the second shortcoming noted above. Losses from a particular type of event (e.g., a windstorm) are likely to generate little updating in the distributions of other event losses (e.g., an earthquake). Thus, separate identification of probability updating and capital market imperfections is more dependable. The cost is that calculating exposures' catastrophe type *and* region entails considerable computational effort. First, the frequency and severity of each type of catastrophe must be estimated *by region*. Then these distributions must be used to derive the distribution of losses on each contract by simulation. The calculation of contract exposure by catastrophe type is taken up in section 5.3 below, after a discussion of the data and the construction of our main variables.

5.2 Data

Our data are built up from four sources. The basic information on catastrophe-reinsurance pricing is provided by Guy Carpenter. Information on the regional market share of insurers is developed from A. M. Best data on insurance premiums written by company. Our estimates of catastrophe frequency and severity are based on Property Claims Services (PCS) data on U.S. catastrophe losses since 1949. Finally, interest rate and CPI data are collected from Ibbotson and Associates and the IMF, respectively.

5.2.1 Guy Carpenter Catastrophe Treaty Data

Our basic data come from Guy Carpenter's proprietary database of catastrophe-reinsurance contracts. Guy Carpenter is by far the largest U.S. catastrophe-reinsurance broker, with a market share of between 30 and 80 percent during our sample period. The contracts brokered by Guy Carpenter cover a variety of natural perils, including earthquake, fire, hurricane, winter storm, and windstorm.

We examine a total of 489 contracts brokered for eighteen national and nineteen regional insurers over the period 1970–94.⁹ The duration of coverage is typically one year. Most contracts have a single mandatory-reinstatement provision.¹⁰ Data on contract inception date, retention (i.e., the retention of the lowermost layer in the contract), limit (i.e., the sum of the limits of all the layers in the contract), and premium (i.e., the sum of the premiums paid for the layers in the contract) are employed. All the contract inception dates are at the start of a quarter.

5.2.2 A. M. Best Market Share Data

To determine the catastrophe exposure of each contract, we must calculate the distribution of contract losses, a random variable for each contract. To do this, we assume that, within each region, each company's exposure is proportional to insurance-industry exposure within the region. We therefore first determine a distribution for insurance-industry losses for each region (by event type) and then multiply this aggregate distribution by an individual insurer's market share to determine the distribution of insurer-specific exposure faced by that company. Using this information, we can calculate the company-specific distribution of losses under each contract.

9. Seven very small regional insurers were dropped from the original Guy Carpenter data. In some of the computations presented below, we focus on a smaller number of national reinsurers, for whom data are available in every year.

10. The reinstatement provision stipulates that, conditional on an event that triggers losses on the contract, the limit is to be mandatorily reinstated (one time only) by the reinsurer after payment of a reinstatement premium by the cedent. It appears that this provision has had only a modest effect on prices, and we ignore its effects. Conversations with brokers suggest that observed prices are approximately 10 percent *lower* than they would have been without the reinstatement premium. This seems surprising (forward contracts are usually priced at zero) but, if anything, leads us to underestimate what premiums would be in the absence of a reinstatement provision.

Our estimates of insurer market shares are developed using data from A. M. Best on insurance premiums written by company, by line of business, by state, and by year. We reduce these multiline market shares to regional catastrophe market shares by applying a modified Kiln formula, which assigns regional weights to premiums in each line of business on the basis of exposure to catastrophes of that line in that region.¹¹ For example, depending on the region, anywhere between 50 and 95 percent of homeowner's premiums are considered as funding catastrophe exposure. The five U.S. regions used for insurer market shares are the Northeast, the Southeast, Texas, the Midwest, and California.¹² We apply this market-share data to all 489 reinsurance contracts selected from the Guy Carpenter treaty database.

5.2.3 Historic Catastrophe-Loss Data from Property Claims Services

As mentioned above, we must determine the distribution of industrywide losses to calculate the catastrophe exposure of each contract. To do this, we estimate the distributions of catastrophe frequency and severity using data from Property Claims Services (PCS). PCS has cataloged all catastrophe losses on an industrywide basis since 1949 by type and U.S. region. The PCS data are widely used as an industry standard.

Prior to estimating the parameters of the frequency and severity distributions, two adjustments are made to the PCS data. First, the losses are converted to 1994 dollars using the CPI. Second, they are modified to take into account shifts in the portfolios of property exposed to loss over the period. A key component of the latter adjustment is the demographic shift toward California, Florida, and Texas that has characterized recent decades. These two adjustments are carried out by Guy Carpenter. Both adjustments are important. Indeed, the second adjustment implies that the same-size event in *real* dollars causes damages that have grown on average by 5 percent per year over the sample period.

5.2.4 Interest Rate and CPI Data

For the purposes of calculating the net present value of payment flows, we use Ibbotson Associates' index of the return on thirty-day U.S. Treasuries. This is collected monthly from 1970:1 to 1995:4. The U.S. CPI is taken from the IMF's *International Financial Statistics*. The frequency is monthly, from 1970:1 to 1995:3.

11. This is a common industry practice. The specific weights used in our Kiln formula are from Guy Carpenter.

12. The regions are constituted as follows: The Northeast comprises Connecticut, Delaware, Maine, Maryland, Massachusetts, New Hampshire, New Jersey, New York, Pennsylvania, Rhode Island, and Vermont. The Southeast comprises Florida, Georgia, Mississippi, North Carolina, South Carolina, Virginia, and West Virginia. Texas comprises Texas. The Midwest comprises Illinois, Indiana, Kentucky, Missouri, and Tennessee. California comprises California.

5.3 Calculation of Exposure and Price

5.3.1 Exposure

In this section, we describe our method of estimating the catastrophe exposure embodied in each excess-of-loss contract. The estimation is carried out in three stages. First, the frequency and severity of each type of event and region are estimated by maximum likelihood for particular families of distributions. Second, a simulated event history is generated by repeatedly drawing from the fitted frequency and severity distributions. Finally, the payouts under each contract in each year of event history are calculated. The mean of the distribution of these payouts is our estimate of the “quantity” of reinsurance, $q_{j,k,t}$, embedded in that particular contract.

The Frequency and Severity of Catastrophes

The first step toward calculating contract exposure is to estimate the frequency and severity of catastrophes using the adjusted PCS loss data. Altogether, there are over eleven hundred catastrophes recorded by PCS. These events are classified into ten categories: earthquake, fire, flood, freeze, hail, hurricane, snowstorm, tornado, thunderstorm, and windstorm.¹³ Many of these events are relatively minor: only 557 have adjusted losses in excess of \$15 million, and only 107 have losses in excess of \$100 million. Four categories of losses are well represented in the set of large losses: earthquake, fire, hurricane, and windstorm.¹⁴ As our primary interest is in exposure to large losses, we confine our attention to these types. Examination of the data reveals that there is some heterogeneity in the losses that arise from windstorms. In particular, a number of the windstorms refer to winter storms (“Nor’easters”) in New England. Accordingly, we split the windstorm category into two subcategories: winter storm, defined to be a windstorm in New England in either the first or the fourth quarter, and windstorm, defined to be all other occurrences of a windstorm.¹⁵

Having defined these five categories of events, we need to make some assumption about regional effects before we can estimate frequency and severity distributions. The simplest assumption would be that, for each catastrophe

13. PCS classifies many events into more than one category. For instance, winter storms in New England, which have on occasion caused substantial damage, are classified first as windstorm and then as hail, freeze, or snowstorm.

14. During the 1949–94 sample period, there were no floods, snowstorms, or thunderstorms with losses in excess of \$100 million. Only one freeze had losses in excess of \$100 million, a \$307 million freeze in Texas in 1989. Three hailstorms and three tornadoes did produce losses in excess of \$100 million, but these are all dated prior to 1970 and so do not appear in our regression analysis below.

15. The assumption that winter storms do not afflict the Midwest may seem strange. The reason is that our regional market-share data are calculated for the Midwest using only five states: Illinois, Indiana, Kentucky, Missouri, Tennessee. The Dakotas, Michigan, Minnesota, Wisconsin, and other characteristically midwestern states are excluded.

Table 5.1 Frequency and Severity Assumptions by Catastrophe Type

Type	Description of PCS data (1)	Assumptions		
		Regions (2)	Frequency ^a (3)	Severity ^a (4)
Earthquake	10 events, all in CA. Frequency appears throughout year.	CA	1: uniform across quarters	1
Fire	19 events: 12 in CA, 2 in MW, 3 in NE, 2 in SE. Frequency higher in quarter 4 and different for CA. Severity comparable across events.	NE, MW, CA	2: CA and NE/SE/MW/TX. Both uniform across quarters	1
Hurricane	48 events: 26 in SE, 22 in NE and TX. Most in third quarter. More severe in Southeast.	NE, SE, TX	8: SE (4 quarterly) and NE/TX (4 quarterly)	2: SE, NE/TX
Winter storm	35 events: in NE in quarters 1 or 4.	NE	1: uniform across quarters 1 and 4	1
Windstorm	352 events: all regions. Frequency differs across regions, but severity is comparable.	NE, SE, TX, MW, CA	20: one for each region and quarter	1

Note: Assumptions for catastrophe frequency and severity distributions are based on catastrophe experience 1949–94. A catastrophe is defined as an event that gives rise to \$15 million or more in insured losses. Column 1 gives a description of catastrophe occurrence by type, 1949–94. NE denotes Northeast, SE Southeast, TX Texas, MW Midwest, and CA California. Columns 2–4 give the assumptions concerning the frequency and severity distributions. The number in the frequency and severity columns represents the number of separately estimated distributions for that type. For example, the number “1” implies that all regions are pooled and that a single, nationwide distribution is estimated.

^aNumber of regional distributions.

type, event occurrences are drawn from a single nationwide frequency distribution while loss sizes are drawn from a single nationwide severity distribution. Given the relative paucity of loss information, this approach helps by pooling the available data. However, the assumption of equal regional distributions is likely to be incorrect. For instance, hurricanes are much less likely to occur in California than in Florida, and the majority of earthquakes occur in California.

As a result, we make specific assumptions regarding frequency and severity on the basis of a careful examination of the 1949–94 catastrophe data. These assumptions are summarized in table 5.1. A catastrophe is defined as an event that gives rise to \$15 million or more in insured losses. Column 1 summarizes the event history for each type. Column 2 reports the regions in which each event type is assumed to occur. Columns 3 and 4 indicate the number of re-

gional frequency and severity distributions estimated for each type. Some of the constraints, such as the assumption that winter storms do not strike California, seem entirely reasonable. Others, such as the assumption that earthquakes do not strike outside California or that winter storms do not hit the Midwest, are less tenable (although see n. 15) and are dictated largely by data availability.

With the assumptions described in table 5.1, there are thirty-three frequency distributions to estimate. We assume that the frequencies are Poisson distributed and estimate the Poisson parameters by maximum likelihood (the estimates are equal to the mean number of events that occur per quarter). Table 5.2 presents the frequency results in four quarterly arrays, by type and region. The estimated frequencies accord with what one might expect. For example, hurricanes are most likely to occur in the third quarter.

We next consider severity. There are six severity distributions, one for each of the catastrophe types identified in table 5.1. We fit two alternative density functions to the empirical severity distribution of each type. The first is a lognormal distribution, with density function for losses l given by $f(l) = \exp\{-[\ln(l) - \mu]^2/2\sigma^2\}/[l\sigma\sqrt{(2\pi)}]$, $l > 0$. The second is a Pareto distribution, with density function $f(l) = \alpha\beta^\alpha/l_{(1+\alpha)}$, $l > \beta$. Once again, the estimation is carried out by maximum likelihood. The fitted distributions are reported in table 5.3. For earthquake, winter storm, and windstorm events, the likelihood ratio test selects the Pareto distribution as the better fit, while, for fire and hurricane events, the lognormal distribution is preferred. However, because the Pareto distribution tends to place a large amount of probability in the righthand tail of the distribution, it does not perform well in attaching reasonable probabilities to large losses. For example, using the estimated Pareto density, the probability that a hurricane in the Southeast generates \$15 billion in losses (given that a hurricane occurs) is almost 10 percent, which appears somewhat high.¹⁶ It might be preferable, therefore, to use the lognormal fit as the baseline severity distribution for all event types. This is the strategy that we adopt. The fitted lognormal distributions are shown in figure 5.9 for losses in the \$0–\$3 billion range.

Simulated Event History

Using these frequency and severity distributions, we are able to simulate an “event history” of catastrophes. From this event history, the distribution of payments under each excess-of-loss contract can be obtained.

Of course, it is not necessary to simulate the distribution of contract payments. In principle, it is possible to determine contract payments analytically. However, analytic solutions are complicated because a contract’s payment is triggered by only a single event, even though that event could be one of five

16. Using PCS data, Cummins, Lewis, and Phillips (chap. 3 in this volume) argue that the Pareto distribution tends to overestimate the probability in the tail of catastrophe severity distributions and that the lognormal fit is to be preferred on these grounds.

Table 5.2 Frequency of Catastrophes, Measured by Their Poisson Parameters, by Quarter, Type, and Region, 1949-94

	NE	SE	TX	MW	CA	NE	SE	TX	MW	CA
	January-March					April-June				
Earthquake					.054					.054
Fire	.031	.031	.031	.031	.125	.031	.031	.031	.031	.125
Hurricane (SE)		.000					.043			
Hurricane (NE/TX)	.000		.000			.033		.033		
Winter storm	.380									
Windstorm		.652	.326	.500	.304	.196	.457	1.109	.935	.000
	July-September					October-December				
Earthquake					.054					.054
Fire	.031	.031	.031	.031	.125	.031	.031	.031	.031	.125
Hurricane (SE)		.370					.130			
Hurricane (NE/TX)	.283		.283			.033		.033		
Winter storm						.380				
Windstorm	.174	.065	.152	.326	.000		.283	.326	.370	.130

Note: Poisson parameter is equivalent to the mean number of catastrophe occurrences per quarter by type and region. If the frequency of each catastrophe type in each region is Poisson distributed $-f(n) = e^{-\lambda} \lambda^n / n!$, where n is the number of events that occur—then the numbers in the table are the maximum likelihood estimates of λ . NE denotes Northeast, SE Southeast, TX Texas, MW Midwest, and CA California. Blank elements of the arrays are 0 by assumption (see table 5.1).

Table 5.3 Fitted Severity Distributions by Catastrophe Type, 1949-94

Distribution and Parameter	Earthquake	Fire	Hurricane (SE)	Hurricane (NE/TX)	Winter Storm	Windstorm
<i>n</i>	10	19	26	22	35	352
Lognormal:						
μ	-2.100	-2.350	-1.233	-1.454	-2.440	-3.039
σ	1.964	1.196	1.610	1.454	1.166	.859
Mean log-likelihood	.006	.752	-.662	-.340	.867	1.772
$\Pr(l > \$5 \text{ billion})\%$	2.915	.046	3.870	1.760	.025	.000
$\Pr(l > \$15 \text{ billion})\%$	.684	.001	.718	.211	.000	.000
Pareto:						
α	.476	.541	.337	.364	.568	.862
β	.015	.015	.015	.015	.015	.015
Mean log-likelihood	.358	.735	-.854	-.556	.875	1.891
$\Pr(l > \$5 \text{ billion})\%$	6.288	4.327	14.110	12.057	3.684	.670
$\Pr(l > \$15 \text{ billion})\%$	3.727	2.389	9.743	8.082	1.973	.260

Note: Results from fitting of lognormal and Pareto distributions to PCS event losses. PCS losses have been adjusted for inflation and population movements by Guy Carpenter. A catastrophe event is defined as one giving rise to insured losses in excess of \$15 million. The density function for the lognormal is $f(l) = \exp\{-[\ln(l) - \mu]^2/2\sigma^2\}/[l\sigma\sqrt{(2\pi)}]$, $l > 0$, while the density function for the Pareto is $f(l) = \alpha\beta^\alpha/l^{(\alpha+1)}$, $l > \beta$. The parameters μ , σ , and α (not β , which is a fixed scale parameter set equal to \$15,000,000) are estimated by maximum likelihood. For a given catastrophe type, estimated mean log likelihoods for the two distributions are comparable and provide a means for choosing between them. The table also shows the probability that an event produces insured losses in excess of \$5 billion and \$15 billion, respectively.




Fig. 5.9 Fitted severity distributions by catastrophe type, 1949–94

different peril types. The single-event clause is in effect a knockout provision, allowing the contract to mature following the first event that generates losses in excess of the retention. For example, it may be that earthquakes are the major large risk for a contract to trigger, but a large freeze in the Northeast in early January could trigger the contract, thereby knocking out the earthquake risk for the remainder of the year.

This knockout provision gives the contract a payment distribution that is very different from that which would apply if the contracts were instead written to cover aggregate losses (i.e., the sum of losses across events). It can also give rise to some paradoxical effects. For example, an increase in the frequency of winter storms may actually *reduce* the total exposure embodied in a single-event contract since it may increase the probability that it matures following a winter storm rather than a devastating hurricane. We look briefly at the value of the knockout provision in the appendix.

We simulate a 1,250-year event history. For each quarter, the following steps are performed: (1) The number of events of each type that occur in each region is randomly drawn from the relevant Poisson frequency distribution (table 5.2). (2) For each event that occurs, a loss amount is randomly drawn from the relevant severity distribution (table 5.3). (3) Finally, all the events that occur in the quarter are randomly sequenced in time. The random sequencing of the events throughout the quarter is an approximation, at best. It is likely, for example, that winter storms occur more frequently in January than March. While it would be preferable to sequence the events on a time scale finer than quarterly, too few events have occurred since 1949 to allow the estimation of this.

Contract Exposure

The exposure of each excess-of-loss contract in our data can be calculated by examining its loss experience in each year of the simulated event history. To take an example, suppose that we are considering a contract brokered by a national insurer with an inception date of 1 April. Let L and R be the contract's limit and retention, and let $m_{j,k}$, $k \in \{\text{NE SE TX MW CA}\}$ be the j th insurer's market share in each of the five regions. The contract's exposure is measured as follows:

1. Split the event history into 1,249 year-long periods measured from 1 April to 31 March.
2. Consider each period in turn. If no event occurs in a period, move to the next period. Otherwise, consider each event in sequence.
 - a) Let the first event be in region k , and let insured losses from this event be l .
 - b) If $m_{j,k}l > R$, the contract is triggered. Measure the reinsurance payment for this period as $\min(L, m_{j,k}l - R)$, and move on to the next period. The contract is no longer in force.

- c) If $m_{j,k}l < R$, no payment takes place, and the contract remains in force. Move on to the next event or to the next period if there are no more events.

This algorithm generates 1,249 observations on the distribution of payments under the contract. The first moment of this distribution is the expected exposure to catastrophe losses. It is easy to derive various conditional loss distributions from the unconditional distribution, such as the distribution of hurricane losses or the distribution of losses from events in the Northeast.

5.3.2 Contract Quantity and Price

We label the expectation of the unconditional distribution $q_{j,t}$, the exposure embodied in company j 's contract at time t . This is interpreted as the quantity of reinsurance purchased. By considering only those reinsurance payments that occur following particular types of catastrophes, it is possible to use the same algorithm to calculate the contract's exposure to each of the catastrophe types listed in table 5.1 above. These exposures by type form the inputs to the calculation of loss weightings in equation (3).

To calculate contract price, we begin with the premium paid for each contract. This is measured simply as the sum of the premiums paid for each layer. Typically, the premiums are paid on a quarterly basis over the duration of the contract. We discount these premium flows back to the contract inception date using the three-month Treasury bill rate. By using the riskless rate, we are equating actuarial present values with true value. Strictly speaking, this assumption holds only under risk-neutral pricing and in the absence of insurer credit risk. However, given that catastrophe losses are uncorrelated with total wealth, risk-neutral pricing is not easily rejected. Furthermore, the use of a risk-adjusted discount rate would, in practical terms, have little import for our results.

Once the net present value (NPV) of the premiums is calculated, it is converted to 1994 dollars using the CPI deflator. Our measure of price is the net present value of premiums divided by contract exposure. Thus, the price of company j 's contract at time t is

$$(6) \quad p_{j,t} = \frac{\text{NPV}(\text{premiums})}{q_{j,t}}.$$

5.4 Graphic Analysis

We are now in a better position to understand the data and computations behind figures 5.2 and 5.3 above. They plot, respectively, the industry quantity and price series quarterly from 1975:1 to 1993:4. The quantity series is the simple sum of exposure across companies. The price series is the exposure-

weighted average of the prices of all contracts in force in that quarter.¹⁷ Note that figure 5.3 shows that risk was sometimes ceded at less than actuarial value (i.e., $p_{j,t} < 1$) during the 1970s and early 1980s.

Hurricane Andrew is responsible for the largest catastrophe loss during our sample period. In the light of this, it is of interest to look at the time series of prices around the time of this event. In particular, we can differentiate between the price-quantity reactions of those contracts heavily exposed to hurricane risk/Southeast risk and those with relatively less exposure. Table 5.4 contrasts the price and quantity responses. From panel A, we see that even those contracts with zero market share in the Southeast show large increases in price in the wake of Andrew. This is supportive of the capital market imperfections view. However, as already discussed, even if there are no other losses occurring at this time, these price responses may be the result of probability updating if the Andrew loss experience revealed new information about hurricane exposure in other regions. Panel B investigates this by sorting contracts according to their hurricane exposure. It turns out that it is those contracts least exposed to hurricane losses that exhibit the largest increase in price. Taken at face value, this suggests that capital market imperfections, and not probability updating, are the most important determinant of the price responses. In order to shed more light on this question, we need to estimate the conditional expectation functions (1) and (2).

5.5 Estimation

5.5.1 Exposure Measured by Regional Market Share

In this subsection, we differentiate events by region and estimate (1) and (2) using regional market shares as proxies for regional exposure. Eight variants of the base specification are estimated, corresponding to different assumptions about the functional form through which losses affect prices. The results are shown in tables 5.5 and 5.6.

Turning to the price regressions first, the β_k are positive and statistically significant in all specifications. The γ_k are positive but generally not very significant. This is *prima facie* evidence that both capital market imperfections and probability updating play some role in determining the response of price to catastrophe losses. However, in all cases, the coefficient on unweighted losses is larger and more statistically precise than that on exposure-weighted losses. This suggests that the supply-side capital market channel is the more dominant of the two.

17. This industry series is based on the contract prices for four insurers that purchased reinsurance through Guy Carpenter in every year from 1975 to 1993. It is representative of the behavior of prices for the other insurers in our database.

Table 5.4 **Event Study of Hurricane Andrew**

	A. Southeast Exposure			B. Hurricane Exposure		
	Mean Exposure	Mean $\Delta \ln(p_{it})$	Mean $\Delta \ln(q_{it})$	Mean Exposure	Mean $\Delta \ln(p_{it})$	Mean $\Delta \ln(q_{it})$
5 most-exposed insurers	.707	.310	.085	.654	.270	-.030
5 least-exposed insurers	.000	.334	-.011	.218	.557	-.138

Note: Comparison of price responses in the year after Hurricane Andrew (20 August 1992–19 August 1993) for different insurers. Panel A contrasts insurers that have high and low exposure to the Southeast (as measured by market share). Panel B contrasts insurers that have high and low exposure to hurricanes. The table shows the mean exposure and the mean price change of the five most extreme contracts in each case. The mean price change for the insurers with lesser exposure to the Southeast is calculated using all fourteen of the insurers that have zero market share in that region.

Table 5.5 The Response of Price and Quantity to Losses, Exposure Measured by Regional Market Shares

	$\sum_{s=1}^2 l_{k,t-s-1}$		$\sum_{s=1}^2 l_{k,t-s-1}^2$		$\sum_{s=1}^4 l_{k,t-s-1}$		$\sum_{s=1}^4 l_{k,t-s-1}^2$	
	$\Delta \ln(p_{jt})$	$\Delta \ln(q_{jt})$	$\Delta \ln(p_{jt})$	$\Delta \ln(q_{jt})$	$\Delta \ln(p_{jt})$	$\Delta \ln(q_{jt})$	$\Delta \ln(p_{jt})$	$\Delta \ln(q_{jt})$
$\sum_k (l_k)$	1.865 (.449)	-.518 (.300)	.103 (.027)	-.028 (.018)	1.736 (.326)	-.284 (.202)	.109 (.020)	-.019 (.011)
$\sum_k w_k \theta(l_k)$	.372 (.261)	-.039 (.160)	.014 (.013)	-.001 (.009)	.519 (.249)	-.169 (.149)	.021 (.013)	-.006 (.008)
R^2	.140	.012	.096	.008	.263	.019	.230	.015
N	435	435	435	435	435	435	435	435

Note: OLS estimates of price- and quantity-response regressions. The dependent variables are 100 times the change in the natural logarithm of contract price and 100 times the change in the natural logarithm of quantity. Losses are measured in billions of dollars. Exposure to each type of regional losses (w) is measured by regional market shares, as in eq. (5). Each column corresponds to a separate regression. Four different assumptions are made about the functional form of distributed lag of losses that affects prices. In cols. 1 and 2, the loss variable is simply the sum of two lagged quarterly losses; in cols. 3 and 4, it is the sum of squared losses from two lagged quarters; in cols. 5 and 6, it is the sum of four lagged quarterly losses; and, in cols. 7 and 8, it is the sum of squared losses from four lagged quarters. OLS standard errors are given in parentheses.

Table 5.6 The Response of Price and Quantity Losses, Exposure Measured by Regional Market Shares

	$\sum_{s=1}^2 l_{k,t-s-1}$		$\sum_{s=1}^2 p_{k,t-s-1}$		$\sum_{s=1}^4 l_{k,t-s-1}$		$\sum_{s=1}^4 p_{k,t-s-1}$	
	$\Delta \ln(p_{j,t})$	$\Delta \ln(q_{j,t})$	$\Delta \ln(p_{j,t})$	$\Delta \ln(q_{j,t})$	$\Delta \ln(p_{j,t})$	$\Delta \ln(q_{j,t})$	$\Delta \ln(p_{j,t})$	$\Delta \ln(q_{j,t})$
$\sum_k \theta(l_k)$	1.759 (.447)	-.468 (.286)	.103 (.028)	-.029 (.018)	1.164 (.345)	-.148 (.212)	.054 (.023)	-.005 (.013)
$w_{NE} \theta(l_{NE})$	6.621 (1.490)	-1.272 (.645)	4.681 (1.228)	-0.835 (.433)	4.514 (1.057)	-.670 (.624)	3.037 (.959)	-.592 (.439)
$w_{SE} \theta(l_{SE})$	.175 (.215)	.116 (.140)	.009 (.012)	.006 (.008)	.463 (.221)	.023 (.118)	.024 (.012)	-.000 (.006)
$w_{TX} \theta(l_{TX})$	3.830 (2.606)	.895 (1.287)	2.622 (2.199)	.404 (.763)	4.595 (2.105)	.241 (1.101)	3.293 (1.833)	-.078 (.590)
$w_{MW} \theta(l_{MW})$	53.052 (37.846)	-45.808 (25.460)	95.048 (110.627)	-135.190 (84.527)	34.090 (30.153)	-14.479 (21.793)	72.853 (66.278)	-12.393 (53.118)
$w_{CA} \theta(l_{CA})$	8.982 (4.599)	-12.595 (4.813)	3.943 (2.444)	-5.589 (2.754)	6.589 (2.571)	-8.817 (2.678)	2.659 (.961)	-2.978 (1.116)
R^2	.220	.048	.197	.035	.306	.060	.283	.047
N	435	435	435	435	435	435	435	435

Note: OLS estimates of price- and quantity-response regressions. The dependent variables are 100 times the change in the natural logarithm of contract price and 100 times the change in the natural logarithm of quantity. Losses are measured in billions of dollars. Exposure to each type of regional losses (w) is measured by regional market shares, as in eq. (5). Each column corresponds to a separate regression. Four different assumptions are made about the functional form of distributed lag of losses that affects prices. In cols. 1 and 2, the loss variable is simply the sum of two lagged quarterly losses; in cols. 3 and 4, it is the sum of squared losses from two lagged quarters; in cols. 5 and 6, it is the sum of four lagged quarterly losses; and, in cols. 7 and 8, it is the sum of squared losses from four lagged quarters. NE denotes Northeast, SE Southeast, TX Texas, MW Midwest, and CA California. OLS standard errors are given in parentheses.

To get a sense of magnitudes, suppose that a \$10 billion event occurs in a particular region. Using the first price specification (which uses two quarters of lagged losses), the \$10 billion loss increases all contract prices in the next year by an average of about 19 percent. Notice that this price increase is independent of contract exposure to the region. Higher exposure to the affected region leads to further, but much smaller, increases in price. A firm that increases its share of the market in the affected region from 0 to 100 percent sees its reinsurance price increase by an average of 4 percent. Thus, much of the increase in price appears to be due to a decline in the supply of reinsurance (or to increases in expected losses *outside* the affected region). Probability updating and capital market imperfections both help, but the capital market effects seem much larger.

Table 5.5 also reports results from equation (2), where quantity reinsured is the dependent variable. The results here are consistent with those presented above: the coefficients on unweighted losses provide strong evidence of capital market imperfections. The δ_k coefficients are negative and statistically less than zero at the 10 percent level in all cases. As we expect from figure 5.1 and figures 5.7 and 5.8, the ϕ_k coefficients on market-share weighted losses are ambiguous in sign.

To calibrate, the first quantity specification suggests that a \$10 billion loss in a particular region leads to a 5.2 percent *decline* in quantity on average over the next year. Meanwhile, for the same size loss, a company that has 100 percent market share in the affected region on average purchases 0.4 percent less reinsurance than if it had a 0 percent market share. Note, however, that this latter effect is neither economically nor statistically significant. In other words, we cannot reject the hypothesis that *relative-exposure* levels have no effect on quantity purchased. But we can reject the hypothesis that purchases do not decline on average subsequent to an event. The quantity results are therefore most consistent with the capital market imperfections story.

In table 5.6, we relax the restriction that all the γ_k and ϕ_k coefficients are equal. This allows the probability-updating effect to differ by region. Indeed, there is some evidence for this in the coefficient estimates, which differ substantially across regions. Here, there is somewhat stronger evidence of probability updating. The cumulative sum of regional losses over the last four quarters appears to affect price positively in all cases. Nonetheless, unweighted losses in each region continue to affect price positively and quantity negatively. All but two of these estimates are statistically significant.

5.5.2 Exposure Measured by Actuarially Expected Loss

The specifications in tables 5.7 and 5.8 are analogous to those in tables 5.5 and 5.6. However, here we use actuarial exposures to weight losses. As discussed earlier, probability updating may occur across, rather than within, regions, and this complicates the identification of capital market effects using regional market shares. By distinguishing losses by type rather than by region,

Table 5.7 **The Response of Price and Quantity to Losses, Exposure Measured Directly by Catastrophe Type**

	$\sum_{s=1}^2 I_{k,t-s-1}$		$\sum_{s=1}^2 I_{k,t-s-1}^2$		$\sum_{s=1}^4 I_{k,t-s-1}$		$\sum_{s=1}^4 I_{k,t-s-1}^2$	
	$\Delta \ln(p_{jt})$	$\Delta \ln(q_{jt})$	$\Delta \ln(p_{jt})$	$\Delta \ln(q_{jt})$	$\Delta \ln(p_{jt})$	$\Delta \ln(q_{jt})$	$\Delta \ln(p_{jt})$	$\Delta \ln(q_{jt})$
$\sum_k \theta(l_k)$	4.143 (1.097)	-1.644 (1.031)	.172 (.056)	-.070 (.055)	3.161 (1.004)	-1.617 (.858)	.180 (.052)	-.081 (.045)
$\sum_k w_k \theta(l_k)$	-.897 (.445)	.505 (.393)	-.033 (.023)	.021 (.021)	-.418 (.448)	.523 (.344)	-.025 (.021)	.024 (.016)
R^2	.136	.018	.101	.013	.238	.023	.222	.022
N	435	435	435	435	435	435	435	435

Note: OLS estimates of price- and quantity-response regressions. The dependent variables are 100 times the change in the natural logarithm of contract price and 100 times the change in the natural logarithm of quantity. Losses are measured in billions of dollars. Exposure to each type of catastrophe loss (w) is calculated directly by simulation. Each column corresponds to a separate regression. Four different assumptions are made about the functional form of distributed lag of losses that affects prices. In cols. 1 and 2, the loss variable is simply the sum of two lagged quarterly losses; in cols. 3 and 4, it is the sum of squared losses from two lagged quarters; in cols. 5 and 6, it is the sum of four lagged quarterly losses; and, in cols. 7 and 8, it is the sum of squared losses from four lagged quarters. OLS standard errors are given in parentheses.

Table 5.8 The Response of Price and Quantity Losses, Exposure Measured Directly by Catastrophe Type

	$\sum_{k=1}^2 l_{k,t-t-1}$		$\sum_{k=1}^2 l_{k,t-s-1}^2$		$\sum_{k=1}^4 l_{k,t-3-1}$		$\sum_{k=1}^4 l_{k,t-1-1}^2$	
	$\Delta \ln(p_{i,t})$	$\Delta \ln(q_{i,t})$	$\Delta \ln(p_{i,t})$	$\Delta \ln(q_{i,t})$	$\Delta \ln(p_{i,t})$	$\Delta \ln(q_{i,t})$	$\Delta \ln(p_{i,t})$	$\Delta \ln(q_{i,t})$
$\sum_i \theta(l_i)$	1.491 (1.137)	-.771 (1.097)	.091 (.055)	-.045 (.053)	.535 (1.126)	.025 (1.060)	-.035 (.089)	-.006 (.073)
$w_{EQ} \theta(l_{EQ})$	65.826 (50.672)	-30.015 (17.103)	44.099 (37.030)	-17.720 (10.130)	24.265 (16.375)	-35.153 (16.931)	26.740 (11.852)	-29.059 (12.808)
$w_{FI} \theta(l_{FI})$	-26.137 (28.439)	-8.489 (20.256)	6.451 (11.162)	-11.843 (6.311)	-36.406 (32.047)	13.440 (21.315)	-3.873 (15.906)	-1.220 (7.531)
$w_{HR} \theta(l_{HR})$	-.218 (.446)	.307 (.410)	-.015 (.022)	.016 (.020)	.266 (.459)	.010 (.410)	.028 (.031)	.004 (.026)
$w_{WS} \theta(l_{WS})$	51.084 (25.298)	13.796 (9.816)	27.597 (20.651)	12.133 (8.235)	21.862 (21.868)	10.452 (12.211)	-5.706 (21.304)	10.512 (14.074)
$w_{WD} \theta(l_{WD})$	9.273 (1.771)	-3.362 (1.200)	5.912 (.762)	-1.883 (.399)	5.905 (1.226)	-2.918 (1.088)	2.726 (.622)	-.881 (.415)
R^2	.249	.038	.264	.032	.301	.056	.291	.039
N	435	435	435	435	435	435	435	435

Note: OLS estimates of price and quantity response regressions. The dependent variables are 100 times the change in the natural logarithm of contract price and 100 times the change in the natural logarithm of quantity. Losses are measured in billions of dollars. Exposure to each type of catastrophe loss (w) is calculated by simulation. Each column corresponds to a separate regression. Four different assumptions are made about the functional form of distributed lag of losses that affects prices. In cols. 1 and 2, the loss variable is simply the sum of two lagged quarterly losses; in cols. 3 and 4, it is the sum of squared losses from two lagged quarters; in cols. 5 and 6, it is the sum of four lagged quarterly losses; and, in cols. 7 and 8, it is the sum of squared losses from four lagged quarters. EQ denotes earthquake, FI fire, HR hurricane, WS winter storm, and WD windstorm. OLS standard errors are given in parentheses.

we can separately identify capital market and probability-updating effects on prices.

The results presented in tables 5.7 and 5.8 are in line with what we found in the previous two tables. First, the unweighted losses in the price equations enter positive and significantly, the coefficients having even larger signs. Specification 1, for example, reveals that a \$10 billion event tends to increase prices by over 40 percent in the following year. Moreover, the aggregated γ coefficient on event-exposure-weighted losses is not positive in the price equations, suggesting that the unweighted losses account for *all* explainable price increases subsequent to events. Second, the unweighted losses in the quantity regressions enter negatively and statistically significantly in three of the four regressions. This is again consistent with capital market imperfections. Indeed, the coefficient is large, suggesting that a \$10 billion event reduces the quantity of reinsurance purchased by between 8 and 16 percent. There is no evidence supporting the presence of event-specific probability updating in table 5.6.

Table 5.7 provides similar results. Unweighted losses influence prices positively and quantities negatively (although not significantly). Squared losses do not seem to have the expected effects, suggesting that large losses may not have proportionately as large effects as smaller losses in the data.

Taken together, these results provide evidence to support the existence of both imperfections in capital supply and probability updating. However, it is the former that accounts for the bulk of price movements in the wake of losses.

5.6 Conclusion

There are at least two candidate explanations for the sharp rise in catastrophe-reinsurance prices and retentions that occurred in the 1990s, a time of unprecedented catastrophe losses. The first holds that capital market imperfections impeded the flow of capital into the reinsurance sector. In the presence of these imperfections, prices are bid up, and quantities fall, owing to the supply contractions that accompany losses. The second explanation is that the changes in prices and quantities have largely been the result of an increase in the perceived frequency or severity of catastrophes. If, after a loss, neither supply nor demand shifts but actuarial probabilities of losses tend to increase, then we would observe price increases and quantity reductions since our quantity measure is derived from time-invariant loss distributions.

To separate out the effects of capital market imperfections and probability updating, we consider two specifications of how loss distributions are updated. First, we assume that probability updating occurs on a region-specific basis and that event losses in a given region may therefore increase perceived future losses within that region but not in other regions. Second, we assume that probability updating occurs within (but not across) event types and that event losses associated with a particular peril may therefore increase perceived future losses from that peril but not losses associated with other perils.

We estimate these specifications using detailed reinsurance contract data from Guy Carpenter. Our findings suggest that, subsequent to losses, price increases and quantity declines are more pervasive across contracts than they should be on the basis of contract-specific exposures to event types and regions. Since cross-sectional variation in exposures should explain changes in prices and quantities but does not, it appears that price shocks are highly correlated across all forms of catastrophe exposure. This lends support to the view that aggregate price and quantity shocks stem from shifts in the supply of capital to the industry. Since reinsurers are financial intermediaries with relatively few fixed factors besides financial capital, the existence of such shifts in supply is evidence of capital market imperfections.

Appendix

The Value of the Knockout Provision in Single-Event Contracts

Using the simulated event history, it is possible to measure the value of the knockout provision that is implicit in single-event contracts. This knockout feature stems from the fact that it is the first event occurrence that triggers the contract.¹⁸ Table 5A.1 illustrates how the knockout feature affects expected payouts under various contract provisions. Columns 1–6 of the top panel consider contracts whose payouts are contingent on the occurrence of particular types of catastrophes. For example, in column 1, the contract can be triggered only by earthquake losses in excess of the retention. Column 8 gives the expected payout on a contract that is structured in the same way as the Guy Carpenter treaties—that is, it pays out on the first event of any type that generates losses in excess of retention. The table shows that the sum of the expected payouts on event-specific contracts (col. 7) exceeds the expected payout on the all-type contract. This is the familiar result that, with imperfectly correlated risks, it is cheaper to buy insurance on a portfolio than it is to buy a portfolio of insurance policies. The bottom panel of the table is structured in the same way, except that it distinguishes events by region rather than by type. Once again, the sum of the expected payouts on the region-specific contracts exceeds that on the all-region contract.

It is noteworthy that, in both cases, the inception date of the contract has only a minor influence on the distribution of payouts.

18. The complexities introduced by the knockout provision form part of the reason why the catastrophe options that trade on the Chicago Board of Trade are written as aggregate rather than single-event contracts.

Table 5A.1 Value of Knockout Provision in Single-Event Excess-of-Loss Contracts

Inception Date	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Quake	Fire	Hurricane (SE)	Hurricane (NE/TX)	Winter Storm	Wind-storm	Sum	All Type
1 January	1.9	2.7	6.5	7.8	1.7	9.8	30.4	22.3
1 July	1.9	2.7	6.5	7.8	1.7	9.9	30.4	22.9
	NE	SE	TX	MW	CA		Sum	All Region
1 January	6.1	8.9	7.8	3.7	3.9		30.4	22.3
1 July	6.3	8.9	7.7	3.8	3.9		30.5	22.9

Note: Expected reinsurance payments (in millions of dollars) for a \$100 million excess-of-\$500-million contract under various contract payment provisions. In the top panel, the contract is single event, but payment is contingent on the type of catastrophe. Thus, in col. 1, the contract matures on the occurrence of the first earthquake event that produces \$500 million in insured losses. Column 7 gives the sum of the expected payments from the contracts in cols. 1–6. Column 8 is the value of expected payments under an all-event contract that matures on the occurrence of any event producing \$500 million in insured losses. The difference between col. 7 and col. 8 is the value of the event knockout provision implicit in single-event, all-event contracts. The bottom panel is structured in exactly the same fashion, except that it distinguishes events by region rather than type. The five regions are Northeast (NE), Southeast (SE), Texas (TX), Midwest (MW), and California (CA).

References

- Cummins, J. D., and P. Danzon. 1991. Price shocks and capital flows in liability insurance. In *Cycles and crises in property/casualty insurance: Causes and implications for public policy*, ed. J. D. Cummins, S. E. Harrington, and R. W. Klein. Kansas City, Mo.: National Association of Insurance Commissioners.
- Cummins, J. D., and J. F. Outreville. 1987. An international analysis of underwriting cycles in property-liability insurance. *Journal of Risk and Insurance* 54:246–62.
- Froot, K. A. 1995. Incentive problems in financial contracting: Impacts on corporate financing, investment, and risk management policies. In *The global financial system*. Cambridge, Mass.: Harvard Business School Press.
- Froot, K. A., and P. G. J. O'Connell. 1997. On the pricing of intermediated risks: Theory and application to catastrophe reinsurance. Working Paper no. 6011. Cambridge, Mass.: National Bureau of Economic Research, April.
- Froot, K. A., and J. Stein. 1998. Risk management, capital budgeting and capital structure policy for financial institutions: An integrated approach. *Journal of Financial Economics* 47 (January): 55–82.
- Gron, A. 1994. Capacity constraints and cycles in property-casualty insurance markets. *Rand Journal of Economics* 25, no. 1:110–27.
- Gron, A., and D. Lucas. 1995. External financing and insurance cycles. Working Paper no. 5229. Cambridge, Mass.: National Bureau of Economic Research.
- Winter, R. A. 1988. The liability crisis and the dynamics of competitive insurance markets. *Yale Journal on Regulation* 5:455–500.

Comment Jeremy C. Stein

“The Pricing of U.S. Catastrophe Reinsurance,” by Kenneth A. Froot and Paul G. J. O’Connell, is an excellent paper. It poses a clear-cut and very interesting question, comes up with an extremely clever empirical approach to tackle the question, and delivers quite compelling results. The question is, To what extent do the dramatic variations in the price of catastrophe reinsurance reflect shifts in reinsurers’ supply curves? In a world of imperfect capital markets, one would naturally expect supply shifts to be associated with large disasters—for example, Hurricane Andrew or the Northridge Earthquake—which badly deplete reinsurers’ existing capital bases.

The question is interesting, not only from the narrow perspective of the catastrophe-insurance market, but to anyone concerned with the economics of financial intermediation. For example, in the context of commercial banking, there has recently been a great deal of interest in the analogous question, Can shocks to banks’ capital influence their willingness to make loans? Although there is by now a large literature that seeks to address this “bank capital

Jeremy C. Stein is the J. C. Penney Professor of Management at the Sloan School of Management, Massachusetts Institute of Technology, and a research associate of the National Bureau of Economic Research.

crunch" issue, it is in many cases plagued by the fundamental identification problem of separating out movements in supply and demand. It is clear that banks lend less when their capital is reduced. It is less clear that this is because of a loan-supply effect—it might instead be that reduced capital is indicative of a deterioration in the lending environment and hence a decline in the demand for new loans. Froot and O'Connell are able to use the unique nature of the catastrophe-reinsurance market to largely get around this type of identification problem and thereby say something more definitive about the role of capital in the supply of intermediation services.

My comments are organized as follows. I will first review the intuition behind the Froot-O'Connell identification scheme. Since I am ultimately quite convinced by their approach, I will only briefly quibble with the details of its implementation. Instead, I will save a little bit of space to discuss some further implications of their results.

The Froot-O'Connell Approach to Identifying Supply Shocks

The Basic Idea

How would one know whether disasters lead to an inward shift in the supply of reinsurance? One obvious thing to do is to look at whether the price of reinsurance rises in the wake of a large disaster. In fact, it typically does, but this is not decisive. There are two potentially confounding factors. First, it may be that the demand for reinsurance also rises after a disaster, as insurers seek more cover. Second, there is the "probability-updating" effect that Froot and O'Connell stress: "price" is measured relative to a static estimate of actuarial value, with the result that, if a big hurricane causes market participants to update their estimates of the actuarial exposure of a given contract, this will incorrectly be interpreted as an increase in "price."

As a next simple step, one might jointly consider information on both prices and quantities. For example, if prices rise and quantities fall after a disaster, this looks more like an inward shift in the supply of reinsurance than an increase in demand. Indeed, this approach is probably all that would be needed if the only confounding factor were demand shifts. But it does not by itself get around the probability-updating problem, as Froot and O'Connell are careful to point out. Because of the way their price and quantity measures are constructed, there can be a "hardwiring" effect where an increase in the perceived probability of a hurricane simultaneously raises the measured price and lowers the measured quantity of reinsurance. Thus, to get convincing identification, one must tackle the probability-updating problem head-on.

This is where the key insight of the paper comes in: while a bad hurricane like Andrew might be expected to change market participants' expectations about the likelihood and costs of future hurricanes, it should not cause similar updating about earthquakes. Thus, if one could somehow document that a bad

hurricane leads to an increase in the price and a decrease in the quantity of earthquake policies, and vice versa, this would be convincing evidence of a supply effect driven by capital constraints. Quite literally, we would have an ideal natural experiment.

Implementation

Unfortunately, the data are not quite as cooperative as one might like. The main problem is that one never actually observes separate prices for earthquake and hurricane policies. Rather, all we know about a given contract is the layer of protection and the regional market shares of the insurer involved. Thus, the exposure of a contract to a given type of peril must be imputed. Froot and O'Connell do this imputation in two ways. The first, simplest method uses only information on insurers' regional market shares. Thus, for example, two insurers who both have 50 percent of their market share in California would be deemed to be equally exposed to earthquakes (since earthquakes are assumed to happen only in California). Obviously, this entails some measurement error since it may be that, in reality, one of the two insurers is more heavily exposed to properties that lie near a fault line. The second method of imputing exposure is substantially more complicated, but it is fair to say that, while it may attenuate this sort of measurement error, it cannot hope to eliminate it.

The measurement error is significant because of the kind of regression specifications that Froot and O'Connell use. For example, in their contract-price equation (1), if the contract exposures (the w 's) are only noisy proxies, this will lead to a downwardly biased estimate of their γ coefficient and most likely an upwardly biased estimate of β . In other words, the method will be biased toward overrejecting the null hypothesis of perfect capital markets. I should say that, as a practical matter, I would be very surprised if measurement error can completely explain away the very strong rejections of the null that show up in their tables 5.5–5.8. Nonetheless, it would be nice if there were some way to demonstrate this point more convincingly.

It should be possible to take some steps in this direction. The idea would be to examine a subset of the data where we know a priori that measurement error is less likely to be an issue. Let me give a concrete example. Suppose we can find some insurers who have literally *zero* market share in California. In this case, we can say fairly confidently that they are completely free of earthquake exposure. (In contrast, once there is some nonzero market share in California, it becomes much harder to say *just how large* the earthquake exposure is.) One could focus in on just these earthquake-free insurers and see how the prices of their contracts respond to an earthquake. If they go up a lot, I would be totally convinced.

It should be noted that this approach is very much in the spirit of what is done in the Hurricane Andrew "event study" in table 5.4, where Froot and O'Connell look at the five insurers with the smallest (albeit still nonzero) expo-

sure to hurricanes and find that their contract prices rose markedly after Andrew. I found this event study to be perhaps the most compelling piece of evidence in the paper, and I would have liked to have seen more in a similar vein.

Implications: What Are the Underlying Capital Market Frictions?

If one accepts Froot and O'Connell's conclusion that the large swings in reinsurance prices seen in recent years are indeed a capital related phenomenon, the natural next question to ask is, What exactly are the underlying primitive frictions that cause reinsurers to become so capital constrained? Without a coherent answer to this question, it would be difficult to evaluate either policy proposals or potential financial innovations that hold the promise of "fixing" the capital scarcity problem.

A simple but important point—one that strikes me as underappreciated in this context—is that capital constraints can have an important effect in steady state only if *two* necessary conditions are both met. First, there must be *flow* costs of adding new external capital. These flow costs might come, for example, from information asymmetries between managers and shareholders that create adverse-selection problems and discourage seasoned equity offerings. In addition, however, there must be *stock* costs of simply holding a large capital buffer on the balance sheet. Were it not for the stock costs of holding a capital buffer, the flow costs of raising new external finance could be rendered irrelevant. In the extreme, a reinsurer could set itself up *ex ante* with a huge capital buffer, so large that, no matter the disaster, it could be assured of never having to seek new external finance under adverse-selection conditions.¹ Thus, the fact that there appear to be persistently binding capital constraints in the reinsurance industry strikes me as indirect evidence that, in addition to any adverse-selection frictions, there must also be important stock costs of holding capital.

Where do the stock costs come from? The two obvious candidates are taxes and agency. On the tax side, with a U.S.-style code, it is obviously inefficient to set up an equity-financed entity that raises billions of dollars from investors and then parks the proceeds in Treasury bills until there is a major disaster—the interest income will be subject to double taxation. Of course, if one takes the view that taxes are the dominant friction, then the ultimate financial innovation is the offshore reinsurance corporation, exemplified by the number of reinsurers that, by setting up in Bermuda, have been able to realize substantial tax advantages.

1. By *ex ante* I mean that it should be possible to raise a large amount of funding to start (and overcapitalize) a new reinsurance company at some initial date before any policies have been written and, hence, before managers have any informational edge over investors. Thus, at this *ex ante* date, there should be no adverse-selection problem in raising the money. Alternatively, a large capital buffer could be accumulated by retaining earnings over a long period of time.

My own guess is that agency considerations are at least as important as taxes in discouraging large buffer stocks of capital.² Put simply, in the standard corporate form, where managers have wide discretion over how to spend internal resources, it can be scary to think about setting them up with a multi-billion-dollar cushion that is not needed for day-to-day operations. The obvious worry is that the money will find a way to get spent, perhaps on writing negative-net-present-value policies in a misguided effort to grow the business.

If agency is indeed the key problem, this suggests thinking about contractual measures that restrict the discretionary uses to which the excess capital can be put. In the extreme, one might think of setting up a “lockbox” arrangement: the money can be used only to pay off the losses on reinsurance contracts that have already been written. But, of course, this extreme lockboxing solution is tantamount to securitization—that is, it completely eliminates managerial discretion and effectively takes the reinsurance out of the corporate form.

While securitization may be helpful in some regards, one should not jump to the conclusion that it can easily “solve” the capital problem. Taking discretion out of the hands of reinsurer management undoubtedly has costs as well as benefits since management presumably is better able to evaluate and price complex risks than would be diffuse capital market investors. Thus, the right question is, *How much* of what is currently done via intermediaries can be usefully taken out of the corporate form? The presence of significant agency problems implies that the corporate form for reinsurance should be used sparingly, only for those particular risks where the ability to exercise managerial discretion is on net beneficial.

2. I do not have any direct evidence to back up this guess. One way to gain some insight might be to study the relative behavior and performance of U.S.-domiciled and Bermuda-based reinsurers over time, as more data become available.