

INVITED ARTICLE

THE EVOLVING MARKET FOR CATASTROPHIC EVENT RISK

Kenneth A. Froot

ABSTRACT

In the last two years, the market for catastrophic event risk has witnessed important change. The first large and truly successful catastrophe ("CAT") bonds have been issued. New exchanges have opened, traded contracts have been created, and indexes of CAT losses have been introduced. The array of products confronting issuers and investors has widened substantially.

This article provides a brief overview of these changes. It also takes a functional approach to diagnosing the problems in the market for CAT event risk in order to understand how future change is likely to occur. Finally, it provides information to companies looking to assess whether these new markets are useful for solving their problems.

THE DISTRIBUTION OF CATASTROPHIC EVENT RISK

The distribution of catastrophic event risk is highly unfavorable without insurance and reinsurance. Individual households and companies are significantly exposed to their own property damage; and because they are not exposed to others' property damage, these entities are not rewarded with a financial gain if no CATs occur. Households and corporations dislike this situation; they face a large potential downside from CATs that affect them yet receive no upside when no CATs occur.

Insurance, reinsurance, and the capital markets provide an obvious redistribution mechanism whereby those who are exposed can protect themselves against CAT loss. Households and companies are risk averse; therefore, they will pay more than actuarial losses for insurance. This demand then translates into an upside for others. When there are no catastrophes, each entity's premiums are redistributed to others in the form of financial gains.

Kenneth Froot is Professor at the Harvard Business School. Ernest Asaff, Christopher McGhee, and Patrick Sullivan of Marsh & McLennan Securities Corp.; Robb McDaniels and Bruce Thomas of IndexCo.; and John Major, Jonathon Norton, and Daniel Bar-Yaacov of Guy Carpenter & Company assisted in the preparation of text and data.


(This article was originally prepared as a special report by Marsh & McLennan Securities Corp. and sponsored by Guy Carpenter & Company, Inc. © 1998 Guy Carpenter & Company, Inc. All rights reserved.)

The Inefficiency of Catastrophic Event Risk Distribution

Clearly, a catastrophe risk distribution mechanism exists. But how well has it functioned historically? The answer is mixed. Many entities, particularly businesses, have large assets that remain greatly exposed to CAT risk. This is because large insurance and reinsurance policies are difficult and sometimes costly to obtain. Indeed, even insurance companies that underwrite policies affected by CAT events have, at times, experienced difficulty transferring substantial quantities of CAT risk to reinsurers.

To demonstrate this, Figure 1 shows the fraction of CAT losses that insurers protect through reinsurance, averaged across insurers.¹ Note that the fraction of protection is high for relatively small CAT losses but low for moderate-sized CAT events. For example, for CAT events generating an \$8 billion industry loss, only about 20 percent of marginal insurer losses would have been covered by reinsurance as of 1994.²

FIGURE 1
Percentage of Marginal Exposure That Insurance Companies Reinsure (by Event Size)



A well-working risk distribution mechanism would have an insurer, like any other company, ceding most of its own CAT exposure. By self-insuring, companies are requiring their equity holders implicitly to provide the insurance.³ After all, equity returns are the first to suffer when there is a CAT event loss.

¹ Figure 1 actually overstates the fraction of losses covered, since it averages across only those insurers that purchase some amount of reinsurance. If insurers that purchase no reinsurance were included, the fraction shown in Figure 1 would decline.

² A CAT event that generates a loss of \$8 billion (1996 dollars) is not extremely unlikely, with a probability of exceedence of approximately 15 percent.

³ Of course, to the extent that the firm has insufficient capital/surplus to sustain a loss, equity holders share the risk with other claimants of the firm: bondholders, trade creditors, and policyholders.

Self-insurance is likely to be inefficient for several reasons. First, equity investors typically expect high average returns. This is largely because equities have large amounts of undiversifiable risk—they move strongly with the entire stock market, and the stock market is volatile. By contrast, catastrophe risks can be fully diversified when embedded in standard portfolios of stocks and bonds.⁴ As a result, both the returns and risks associated with CAT exposures are likely to be lower than those associated with equities. Equity investors would prefer a pure high-risk/high-return investment and would rather not have the “reward dilution” associated with CAT exposures.

Second, equity investors typically expect a long-lived, complex risk. (Even if their trading horizons are short, equity investors know that the current stock price telescopes back both short-horizon as well as long-horizon perceptions about the firm’s prospects.) Equity risks are therefore not targeted to precise events over precise periods of time. Rather, they tend to be fuzzier and evolve over time. Their precise nature is, at least partially, opaque.

CAT risks, by contrast, can be analyzed as event risks. Writers of CAT reinsurance receive a predetermined premium in exchange for protection against a pre-specified event, with a predetermined maximum loss over a predetermined period of time. Whether embedded in a traditional reinsurance contract or in a cutting-edge CAT bond, the event’s nature is clearly defined. Even though the probability of the event may be uncertain, the risk does not evolve much and the scenarios that produce losses and gains are fairly transparent. The event risk is studied more when it has to stand alone. Modeling firms attempt to quantify the likelihood and severity of dollar damages, and then specialists and investors review the models. Much more light is cast on the risk than if it remained unquantified and amorphous, embedded in equity.

The clarity and specificity of event-risk exposures make it easy to separate CAT risks from other risks embedded in equity, such as management risk, firm-specific product risk, firm-specific cost risk, industry-specific product risk, industry-specific cost risk, and general macroeconomic risk. It is fundamentally inefficient to take a clearly defined event risk and embed it into an ambiguous, multipurpose instrument like equity.

Closed-end mutual funds are a good example of the costs of clouding up clearly defined risks. Closed-end funds invest in publicly traded securities and then sell stakes in their portfolio to shareholders, much as open-end mutual funds do. The difference is that open-end funds allow shareholders to sell their shares back to the fund at a price dictated by the net asset value of the portfolio. Closed-end funds do not automatically buy and sell their shares. Rather, a shareholder wishing to sell must find another investor who wishes to purchase the shares. The price of the closed-end fund shares, like the price of most traded stocks, must find its own value in the marketplace in accordance with supply and demand.

There is, however, a puzzle associated with closed-end fund shares. Their prices are, on average, considerably below the net asset values of the shares they own. This cannot happen with open-end fund shares. Closed-end share discounts average about 10 per-

⁴ For discussion regarding the lack of correlation of CAT risks with traditional asset classes, see Kenneth A. Froot, Brian Murphy, Aaron Stern, and Steven Usher, “The Emerging Asset Class: Insurance Risk,” Guy Carpenter & Co., Special Report, July 1995.

cent to 20 percent and are pervasive across funds. Why is there such a discount? The reason is that shareholders cannot observe or control what managers do next with the portfolio. If managers make bad trades, an investor would prefer to liquidate the portfolio at its net asset value. Investors will pay a lower price for a manager-controlled portfolio of stocks—that is, they demand a higher return on their capital when its control is given to a manager. The lack of transparency, clarity, and control associated with managerial discretion creates these higher capital costs. Closed-end funds are therefore expected to pay average returns in excess of what the underlying assets earn.

In this example, the underlying stocks in closed-end fund portfolios are analogous to clearly defined CAT risks: the risks they represent are transparent and specific. Closed-end fund equity is analogous to other companies' equity. Just as closed-end funds could raise their market value by selling off a portion of their portfolios and returning the cash to investors, companies with CAT risk could maximize their value by taking the CAT risk out of their equity and letting it trade on a separate market for CAT event risk.

Not incidentally, a direct implication of the closed-end fund argument is that reinsurers may be inefficiently funded so that their reinsurance is expensive. Reinsurers underwrite pools of CAT risk, retaining most and effectively self-insuring through their equity holders. Pooling helps diversify the CAT risk, but it does not give equity investors the high-octane risk and reward they seek, nor does it retain the clarity of the component event risks. Reinsurers who are currently struggling with the need to earn "equity-like" returns on their underwriting activities increasingly have too much equity and too little event risk protection to be fully efficient.

Reasons for Inefficient Distribution of Catastrophe Risk

The above arguments suggest that CAT event risk sharing is currently inefficient. Why is this the present state of affairs? Reasons include the following.

Lack of Objective Information Technology and computer modeling of natural perils have only recently reached the point where the risks can be cheaply and objectively clarified. If little can be said objectively about the risk of an event, then risk sharing is less likely to occur, because buyers (i.e., assumers) of the risk worry that they will be sheep and that sellers will be wolves. CAT risk tends to accumulate and firms self-insure.

High Cost A second but related reason for the inefficient distribution of CAT risk is that it is costly to back CAT protection with capital. These "frictional" costs can add up to be a significant portion of the premium. For example, suppose that financial technology is such that the cost of assembling a pool of funds from investors is 2 percent of the size of the pool. These "deadweight" costs of creating the pool can be large relative to the premium that the pool can be used to generate. Thus, these costs make it expensive to provide protection for unlikely events. Of course, as financial technology improves so that it is cheaper to assemble the pool, the inefficiency diminishes.

Barrier to Entry Caused by Lack of Objective Information Another reason that the institutional arrangement may be inefficient is that the lack of objective information acts as a kind of barrier to entry. When objective information is costly to assemble, a greater investment is required to get into the underwriting business. Indeed, when objective

information is in short supply, markets tend to be organized around relationships and reputation. By contrast, when objective information is plentiful, markets tend to be organized around transactions, with the players being more interchangeable. Because newcomers are discouraged from entering the market, the incumbents who specialize in underwriting CAT risks (e.g., CAT risk reinsurers) can more easily charge high prices.

Competition and innovation are making today's markets more efficient. It is therefore useful to understand where this greater efficiency will take us. As we have argued, we are likely to see that companies will maximize their market values best by ceding CAT event risks into a CAT event risk market and nonspecific risks into the equity market. As Adam Smith saw centuries ago, specialization is the key to greater efficiency.

Evidence of Improving Efficiency

CAT risk distribution has been, and to a degree still is, inefficient. But the real emphasis here is on the forces of change. The recent increase in efficiency is, of course, hard to measure. However, considerable evidence exists that is consistent with recent improvement.

Most important is the evidence that the price of CAT event protection has fallen for the past four years. Figure 2 shows the average rate-on-line (ROL) of catastrophe excess-of-loss contracts for the decade from 1989 to 1998.⁵ ROL increased substantially following Hurricane Andrew in 1992. By 1993, ROLs averaged 2.5 times the 1989 level. Then, beginning in 1993, rates began to fall. ROL fell steadily from 1994 through 1998. As indicated in Figure 2, the ROL of today's market is at approximately the same level as in the late 1980s—before Hurricane Andrew and the Northridge earthquake.


However, ROL can be a misleading measure of the price of reinsurance. For example, suppose ROL stays constant, but the retention (i.e., deductible) specified in the reinsurance contract falls. All else being equal, this means that the reinsurance is invoked more frequently. In effect, more protection has been purchased for the same overall cost, so the price per unit has actually fallen.

In addition to showing ROL, Figure 2 provides an improved measure of price. The price index, compiled and calculated by Guy Carpenter & Company, Inc., adjusts ROL for changes in retention. Clearly, the price index shows an even more dramatic change in price, up in the immediate aftermath of Hurricane Andrew, then down thereafter. Thus, the basic impression that comes out of Figure 2 is not due to mismeasurement—prices did indeed fluctuate substantially in the 1990s, as it first appears.

Capacity fluctuations are one explanation for these price changes. Prices rose in the early 1990s because capacity was reduced (i.e., the supply of reinsurance fell as a result of catastrophe losses from events). Subsequently, there were few large CATs, which led to a regeneration of capacity. Of course, aficionados will point out this "supply-side" story is not the only possibility. There is a second explanation: price changes might also have been demand driven. That is, demand for reinsurance increased following Andrew and Northridge and then fell with the paucity of additional CATs. It is not imme-

⁵ Rate-on-Line is the ratio of premium to limit (i.e., maximum possible loss covered by the contract), averaged across contracts. Data have been compiled by Guy Carpenter & Company, Inc. using a sample of reinsurance contracts it brokered.

FIGURE 2
Price Level of Reinsurance Contracts (1989-1998)



diately clear which force—supply or demand—was more important in generating the prices shown in Figure 2.

However, the two explanations can be distinguished by examining the two curves in Figure 2 more carefully. The fact that the price index amplifies the movements in ROL implies that less reinsurance was purchased as prices rose and more was purchased later when prices fell. This pattern is not consistent with a demand effect, which would say that more (less) reinsurance would be purchased during times of increasing (decreasing) prices. Only supply shifts could generate the price and quantity patterns observed here.

In other words, over the last four years, the price of reinsurance has fallen, while the quantity of reinsurance purchased has risen. (The quantity purchased has increased particularly strongly in the last year of the sample. This suggests that for the years immediately preceding 1997, prices declined less rapidly than suggested by the ROL.) Only changes in supply could account for this pattern. And increases in supply come from greater competition and efficiency in production. In some sense, the reinsurance market is like the computer chip market—increasing efficiency allows prices to decrease even while quantities soar.

There is also evidence that the improved efficiency in the provision of reinsurance has affected the reinsurance buying patterns of insurers. Figure 1 shows that over time insurers have come to purchase greater amounts of coverage for mid-sized and large events. Beginning in 1994, the fraction of large losses being reinsured may be low, but it is nevertheless much higher than before. Greater competition in the reinsurance market seems to have improved the sharing of CAT event risk.

FINANCIAL INNOVATIONS

Financial innovation permits greater competition among capital suppliers. In the last few years, clear and well-defined CAT event risks have been brought directly to investors for the first time. This direct sourcing of investor capital can improve efficiency by removing layers of management, separating CAT risk from equity risk, and improving the information with which those who bear the CAT risk have to work. In this section, we discuss several of the mechanisms that have been used thus far to securitize CAT event risk.

Table 1 provides a listing of most of the important securitized CAT event risk transactions to date. Because these are the first transactions to occur, there is as yet little standardization. Each transaction is different and challenging in its own way. The purpose of this section, however, is not to uncover the fine details of each. Examples are used from Table 1, but the goal is to convey a broad characterization of the securities and the way they work. With an emphasis on the functional characteristics of these innovations, it is possible to think about the opportunities and limitations for new, yet-to-be-completed transactions.

There are two basic methods of delivering CAT event risk directly to investors: financing transactions and risk-transfer transactions. Financing transactions are those that, upon the occurrence of an event, exchange funds from investors with securities. In a financing transaction, the securities require the insured to repay investors over time, and the exchange of securities for cash is done at fair value at the time of the exchange. Thus, when the CAT event occurs, the investor loses no money: \$100 worth of securities is received in return for \$100 in cash. The main risk faced by investors, therefore, is that they are obliged to come up with the cash quickly. A standard credit facility with a bank is an example of a pure financing transaction (however, drawdown on a credit facility is generally not tied to CAT events).

Risk-transfer transactions are different. Here, the investor loses when the CAT event occurs. Specifically, the investor purchases securities with predetermined event-linked payments. This exchange occurs before the CAT event. If and when a CAT event occurs, investors receive less than full or no repayment. In this way, investors share in the event-linked loss of the insured. If no event occurs, investors share the insured's upside. With risk transfer, the insured company and its stockholders bear less CAT event risk than they otherwise would. Reinsurance and insurance are examples of more or less risk-transfer arrangements, though they are not capital market instruments and they sometimes contain financing elements.

Risk-transfer solutions are likely to be a richer area for innovation than risk-financing transactions. Better risk sharing—which is critical—can be accomplished only through better risk transfer. On the other hand, contingent risk financing can be accomplished in a number of ways. For example, the ability to perform straight financing operations (through debt, equity, or any other instrument) whenever desired obviates any strong need for contingent financing. Moreover, financing methods for general-purpose contingencies—such as bank credit facilities—are close substitutes and are widely available.

TABLE 1
Overview of Major Catastrophe Securitization Issues, 1995-1998

<i>Issuer</i>	<i>Date Closed</i>	<i>Amount (\$mm)⁵</i>	<i>Maturity</i>	<i>Trigger Type</i>	<i>Coupon</i>	<i>Probability of Attachment</i>
Normandy Re ¹	prior to 95	25+	3 yr	PCS Index		1%
ACE Ltd. ¹	1995/96	45	14 mo	PCS Index	LIBOR+550	2%
CAT Ltd. ¹	1995/96	50	5 mo	UNL	UST+450	1.1%
CA Earthquake Authority ²	1995/96	1000	4 yr	UNL	LIBOR+1075	1.27%
AIG, PX Re	5/96	10	20 mo	Index	0	25%
Georgetown Re (notes), St. Paul Re	12/96	44.5	11 yr	UNL	6.097%	Various
Georgetown Re (pref. shares), St. Paul Re	12/96	24	3 yr	UNL	14.15%	Various
Hannover Re II	12/96	100	5 yr	UNL	50% profit	20%
Winterthur	1/97	SF399.5	3 yr	Index	2.25% SF	20%
SLF Re, Reliance National	4/97	10	16 mo	Index	LIBOR+900	15%
USAA, Residential Re A-1	6/97	163.8	12 mo	UNL	LIBOR+273	0.7% ⁴
USAA, Residential Re A-2	6/97	313.2	12 mo	UNL	LIBOR+576	0.7% ⁴
Reinsurer Swap	7/97	35	11 mo	Index	LIBOR+519	0.7% ⁴
Swiss Re A-1	7/97	42	2 yr	PCS Index	LIBOR+255	0.41% ⁴
Swiss Re A-2	7/97	20	2 yr	PCS Index	8.645%	0.41% ⁴
Swiss Re B	7/97	60.3	2 yr	PCS Index	10.493%	0.68% ⁴
Swiss Re C	7/97	14.7	2 yr	PCS Index	11.952%	Approx. 1%
Parametric Re -Notes	11/97	80	10 yr	Parametric	LIBOR + 430	0.70% ⁴
Parametric Re -Units	11/97	10	10 yr	Parametric	LIBOR + 206	0.70% ⁴
Trinity Re A-1	2/98	22.036	9 mo	UNL	LIBOR + 182	0.83% ⁴
Trintity Re A-2	2/98	61.533	9 mo	UNL	LIBOR + 436	0.83% ⁴
USAA, Residential Re	6/98	450	11.5 mo	UNL	LIBOR + 416	0.60% ⁴
Yasuda, Pacific Re	6/98	80	5-7 yr	UNL	LIBOR + 370 ³	0.94%
Reliance National II	6/98	20	9 mo	Index	LIBOR + 925	11.06% ⁴
USF&G, Mosaic Re	7/98	21	1 yr	UNL	LIBOR + 820	2.75% ⁴
Tranche 1—Class B Notes						
USF&G, Mosaic Re	7/98	24	1 yr	UNL	LIBOR + 440	0.61% ⁴
Tranche 2—Class A Notes & Units						LIBOR + 210

Sources: Goldman Sachs, Sedgwick Lane Financial, Guy Carpenter & Company, Inc., and Morgan Stanley & Co.

¹ Deal not closed—data represent proposed plan.

² Deal withdrawn and replaced with a traditional reinsurance contract closed with Berkshire Hathaway.

³ The Pacific Re terms also included a second-event drop down mechanism which pays LIBOR + 950 bps and has a 5.12% attachment probability.

⁴ Percentages represent expected loss.

⁵ Numbers represent net proceeds (as opposed to risk capital).

Financing Catastrophic Event Risk: Contingent Surplus Notes

In terms of financing linked to CAT events, there are two main prototypes. The first is an event-linked issuance of debt. Insurers who have done this have typically issued surplus notes, which are basically debt instruments that for regulatory purposes may be treated as capital. In spite of their status as statutory surplus, surplus notes generally oblige the issuer to repay the funds on a fixed schedule. As a result, these transactions cannot transfer risk; CAT event losses are still borne by the company and its equity holders. The contingent part of contingent surplus notes is that the investors agree to accept the surplus notes not when the deal is struck, but when a particular event, such as a CAT event, occurs.

Table 2 lists surplus note issues. Most of these issues are straight surplus notes: insurers borrowing funds directly, which they then use as statutory capital. Investors receive the repayment in the form of interest and principal, or if the issue is structured as a preferred stock, in the form of dividends and principal.

Only a few companies have issued contingent surplus notes: Nationwide, Hannover Re, and Arkwright. In these transactions, investors place funds into a safe account and receive account-ownership certificates. Initially, the account holds U.S. Treasury bills. Interest on the Treasuries passes through to investors, so it is as though investors are holding safe government securities. However, there is an additional wrinkle: under a set of pre-established conditions, the insurer (or reinsurer) can withdraw the liquid Treasury securities from the account and replace them with its own surplus notes. At that point, investors would receive interest and principal from the issuer, so they would bear the credit risk associated with potential issuer default. Investors are compensated in two ways for this added risk. First, the surplus notes pay a higher coupon than U.S. Treasuries; second, starting at the time the safe account is established, the issuer pays investors a premium for the right to “put” its own surplus notes into the account and withdraw the funds. In a pure financing transaction, the higher coupon on the surplus note would by itself be sufficient to compensate investors for the possibility of issuer default. In that case, the “put” premium would be compensation only for liquidity—that is, for knowing that the facility is there when it is needed.

Financing Catastrophic Event Risk: Contingent Equity

An alternative to contingent debt is, naturally, contingent equity. The setup is similar: the insurer pays a fee to investors for the right to issue equity later in return for cash. When issued, the equity would fetch the going market price, so that investors would receive \$100 worth of stock for \$100 of cash. What makes contingent equity different from a vanilla equity issue is the modifier “contingent”: the trigger for the issue is CAT-related event losses incurred. As long as investors receive the stock issue at the then-fair price, contingent equity is a pure financing vehicle. Although investors do not lose money when the equity is “put,” they require nevertheless a payment in return for the contingency. This payment can be thought of as compensation for the investor’s remaining sufficiently liquid to purchase the put equity on demand.

In practice, the transactions involving contingent equity have been slightly more complex. Partly as a result, they may include elements of risk transfer as well as risk financing. A good example of a contingent equity transaction was a \$100 million issue by a Bermuda reinsurer. The reinsurer paid approximately \$2.4 million per year to investors

TABLE 2
A Sample of Large Surplus Note Issues

<i>Issuer</i>	<i>Issue Date</i>	<i>Maturity (years)</i>	<i>Principal (\$mm)</i>	<i>New Issue Spread To Comparable</i>	
				<i>Offer Yield</i>	<i>Treasury (bp)</i>
Anthem	3/26/97	30	200.0	9.04 %	204
Jackson National	3/13/97	30	250.0	8.18	120
Lumbermens Mutual	6/19/96	30	400.0	9.18	206
Equitable	12/13/95	10	400.0	6.98	124
Equitable	12/13/95	20	200.0	7.73	181*
Metropolitan Life	11/08/95	30	250.0	7.00	75
Metropolitan Life	11/08/95	20	200.0	7.73	143
Minnesota Mutual	9/21/95	30	125.0	8.25	169
Prudential	7/17/95	13	100.0	n.a.	
Prudential	6/29/95	30	250.0	n.a.	
Prudential	6/29/95	12	350.0	n.a.	
Liberty Mutual	5/11/95	30	150.0	8.52	153
Liberty Mutual	4/27/95	12	250.0	8.24	117*
Nationwide (Contingent)	2/13/95	30-10	392.0	n.a.	220
Farmers	7/25/94	10	100.0	n.a.	
Farmers	5/11/94	30	300.0	n.a.	
Principal Mutual	3/03/94	50	100.0	8.09	125*
Principal Mutual	3/03/94	30	200.0	7.94	110
John Hancock	2/25/94	30	450.0	7.44	71
Massachusetts Mutual	2/22/94	30	300.0	n.a.	
Nationwide	2/14/94	30	300.0	n.a.	103
Nationwide	2/14/94	10	200.0	n.a.	68
New England Mutual	2/03/94	30	150.0	7.94	163
General Amer. Life	1/14/94	30	107.0	7.72	142
New York Life	12/08/93	30	300.0	7.52	120
New York Life	12/08/93	10	150.0	6.40	73
Metropolitan Life	10/28/93	30	300.0	n.a.	130
Metropolitan Life	10/28/93	10	400.0	n.a.	95
Prudential	4/21/93	10	300.0	n.a.	112

*Note: Spreads for 12-yr and 20-yr issues are calculated from interpolated yields (straight line) using 10-yr and 30-yr Treasuries.

50-yr issue is shown as spread to 30-yr Treasury.

Source: Securities Data Corporation and USAA

for the right to sell them \$100 million in preferred shares. The reinsurer was entitled to exercise this right if its losses from a single CAT event exceeded \$200 million or if its aggregate losses for any year exceeded \$250 million. The preferred shares paid dividends of LIBOR plus a spread. The spread was set according to a predetermined schedule—higher if the reinsurer’s credit rating was low at the time of the issue and lower if its credit rating was high at the time of the issue. This helped insure that investors would suffer only minimally if CAT losses reduced the reinsurer’s creditworthiness. Once issued, the preferred shares were convertible to the reinsurer’s common equity. The conversion was to be effected at the 30-day average stock price prior to conversion.

In order to be predominantly a financing vehicle, this structure attempts to minimize any risk transfer component. The dividends on the preferred stock are set according to the credit rating at the time of the contingent issue. And the preferred stock is converted to common at the current fair price. Thus, while CAT losses may trigger the reinsurer’s need for financing, they are unlikely to impose a loss on investors. Just as in the contingent surplus note transactions, investors need only supply the cash to buy fairly priced securities in the aftermath of CAT event losses.

Thus far, there have been relatively few contingent risk-financing transactions—debt or equity—even in the presence of a large market in finite reinsurance. There are apparently too many close substitutes available at competitive cost. Reliable financing in the aftermath of a CAT event can be found either through an on-the-spot ex post issue of a debt or equity or through a prearranged ex ante contract that allows drawdown at any time. Indeed, the distinction between even the ex post and ex ante solutions is blurred by conventions such as shelf registration, which allows for very rapid access to commercial paper or other medium funds.

Transferring Catastrophic Event Risk: Act-of-God Bonds

In most instances, however, financing CAT event risk is not enough. Suppose, for example, a CAT loss occurs that is larger than the going-concern value of an insurer. Could the insurer arrange a pure financing transaction in anticipation of that loss? The answer is no. Investors would not provide contingent financing because their securities would be dead on arrival, worth less than their purchase price. The going-concern value would be lost if the insurer had to rely only on risk financing. Risk transfer, however, can preserve the going-concern value. By ensuring that its CAT event exposures are shared, the firm can safeguard its future.

Standard insurance and reinsurance contracts are reasonably elegant techniques for transfer of CAT risk. But like any real-world technique, they are not perfect. First and foremost, the cedent bears the credit risk of the insurance or reinsurance counterparty. There is usually no ironclad “quadruple A” guarantee that the insurance protection will be there when the insured becomes entitled to it. The problem can be diminished—though not eliminated—by buying reinsurance from a collection of reinsurers. Moreover, dividing the transaction up into pieces at times can be costly.

Second, in the past, CAT reinsurance agreements included the tacit understanding that there was a financing element to the transaction. The traditional reinsurance transaction was in fact not pure risk transfer, and this was a shortcoming. After a cedent declared a loss under the contract and the loss was paid, it was assumed that the cedent

would continue buying reinsurance from the same provider and that it would voluntarily do so at a higher price. The higher price didn't so much reflect an increase in the probability of loss as it provided the reinsurer the opportunity to recoup paid losses. *Pure risk transfer transactions*, by contrast, are "memoryless"—their terms are set according to the risk at hand and are otherwise independent of previous outcomes. Of course, there is nothing about the reinsurance contract itself that requires consideration be paid for past claims. So this "shortcoming" is more a matter of industry culture. And, perhaps as a consequence, this aspect of reinsurance is likely to change going forward.

CAT bonds perform a risk transfer function similar to that of insurance and reinsurance. They also avoid the shortcomings mentioned above. With a CAT bond, investors place cash into a safe account from which they receive interest. This account is normally structured as a safe trust. The proceeds of the account are available to the cedent only if a CAT event loss occurs. In return for this benefit, the cedent pays into the account a premium, which is transferred (in addition to interest) directly to investors. If there is no event loss, the safe account remains just that—safe. After interest and premium is paid to investors for a pre-specified period of time (usually from one to three years), the full principal is returned to investors. However, should there be an event loss, the investor will receive only that portion of principal and interest that is left after paying the cedent's event losses.


The first thing to note about this structure is that it resolves some of the shortcomings of traditional insurance and reinsurance contracts. The cedent bears no credit risk. The funds are remote from any other claims or possible uses while in the safe account; they are there for the sole purpose of paying cedent event losses in case there is an event. Also, because the CAT bond is purchased in the capital markets, it is a memoryless transaction with no judgments about past events. There is no relationship to manage between investors and the insured. Investors cannot extract a "repayment" for the losses of prior CAT bond investors. Prior losses may raise the perception of future losses, but in no sense would the cedent in a transaction be paying back the investors who lost last time around.

The second thing to note about CAT bonds is that they may have an unexpected and surprisingly widespread positive side effect: improving the general clarity of corporate bond analyses and ratings. With CAT models serving as reference points, a given CAT bond will have a certain likelihood of paying in full and an average recovery payment when payment isn't full. The promised yield on the bond will depend on both. That is, investors require a higher yield to compensate either for a lower probability of full repayment or for a lower recovery amount, given default. The rating, too, will need to take both probability and recovery into account. Figure 3 shows how Fitch summarizes the relationship between rating, probability, and recovery. Each rating curve in the exhibit is generated by a set of probabilities and recovery rates. The higher the probability of full repayment, the lower the recovery rate must be if the same rating level is to be maintained.⁶

In the past, rating agencies have never been precise about the exact combinations of probability and recovery that result in a given rating for any kind of bond. In fact, some

⁶ We thank David Mordecai of Fitch/IBCA for providing this figure and the underlying analysis.

FIGURE 3
Fitch Iso-Rating Curves



Source: Fitch/IBCA

rating agencies (e.g., Standard & Poor's) have maintained that they evaluate only probabilities and not net recovery when assigning ratings. In the context of CAT, this makes little sense. And, when one thinks about it, this makes little sense in any context. Recovery rates clearly affect investor risk and reward, so they should also affect ratings. It seems that precision modeling of CAT events (however rudimentary it may be) has generated additional clarity on what ratings mean. Over time, this will lead to improvements in the transparency of the ratings process.

The third thing to note about in the CAT bond approach is that it, too, has shortcomings. Clearly, it is costly to set aside enough in an account to pay the largest possible loss under the cedent's contract, no matter how unlikely that loss is. Consider the fact that approximately \$186 billion in capital and surplus backs all worldwide property-liability reinsurance contracts issued by the world's top 100 reinsurers.⁷ If these tradi-

⁷ Figure for 1996 from Standard & Poor's.

tional contracts were alternatively structured as CAT bonds, they would require the sum of all limits—nearly one trillion in capital. While it is not unthinkable that such a large sum of money could be raised, it is certainly a very large amount—on the order of the amount tied up in securitized mortgages. It would seem inefficient to raise the maximum potential loss for each and every risk, especially those risks for which the probability of loss is tiny. To the extent that the losses are unlikely and uncorrelated, it becomes more efficient to pool the risks together in a fund. That looks more like the traditional solution, reinsurance.

To demonstrate in greater detail how a CAT bond works, it is useful to examine the 1997 USAA transaction, the largest single securitized risk transfer to date.⁸ In this case, the safe account was a trust, administered by an offshore special purpose reinsurer (SPR) called “Residential Re.” The SPR had no business purpose other than to sell a one-year \$400 million reinsurance contract to USAA and to issue \$400 million in risk-transfer securities to collateralize fully that reinsurance.⁹ The proceeds of the issuance were held in a trust and invested in highly rated, short-term investments such as commercial paper. In the event of a catastrophe, the trustee would have sold the investments to cover 80 percent of USAA’s losses in excess of \$1 billion (until the \$400 million is exhausted). In return for this reinsurance, Residential Re received a premium from USAA of 600 basis points (\$24 million). The premium, along with virtually all of the interest on the commercial paper, went directly to investors, regardless of whether USAA experienced a loss.

To fund the reinsurance, Residential Re issued securities of two types: principal variable and principal protected. If there was a loss, principal variable investors would have lost some or all of their initial investment. These notes paid interest at a rate of LIBOR plus 575 basis points (this is essentially interest plus the reinsurance premium of 600 basis points, less about 25 basis points for costs).¹⁰ The principal protected securities would have had their principal repayment delayed for ten years in the event of a loss, with a reduction in interest along the way. The principal protected securities paid LIBOR plus 273 basis points. Figure 4 and Table 3 provide additional details of the transaction structure.

Transferring Catastrophic Event Risk: Standardized Contracts

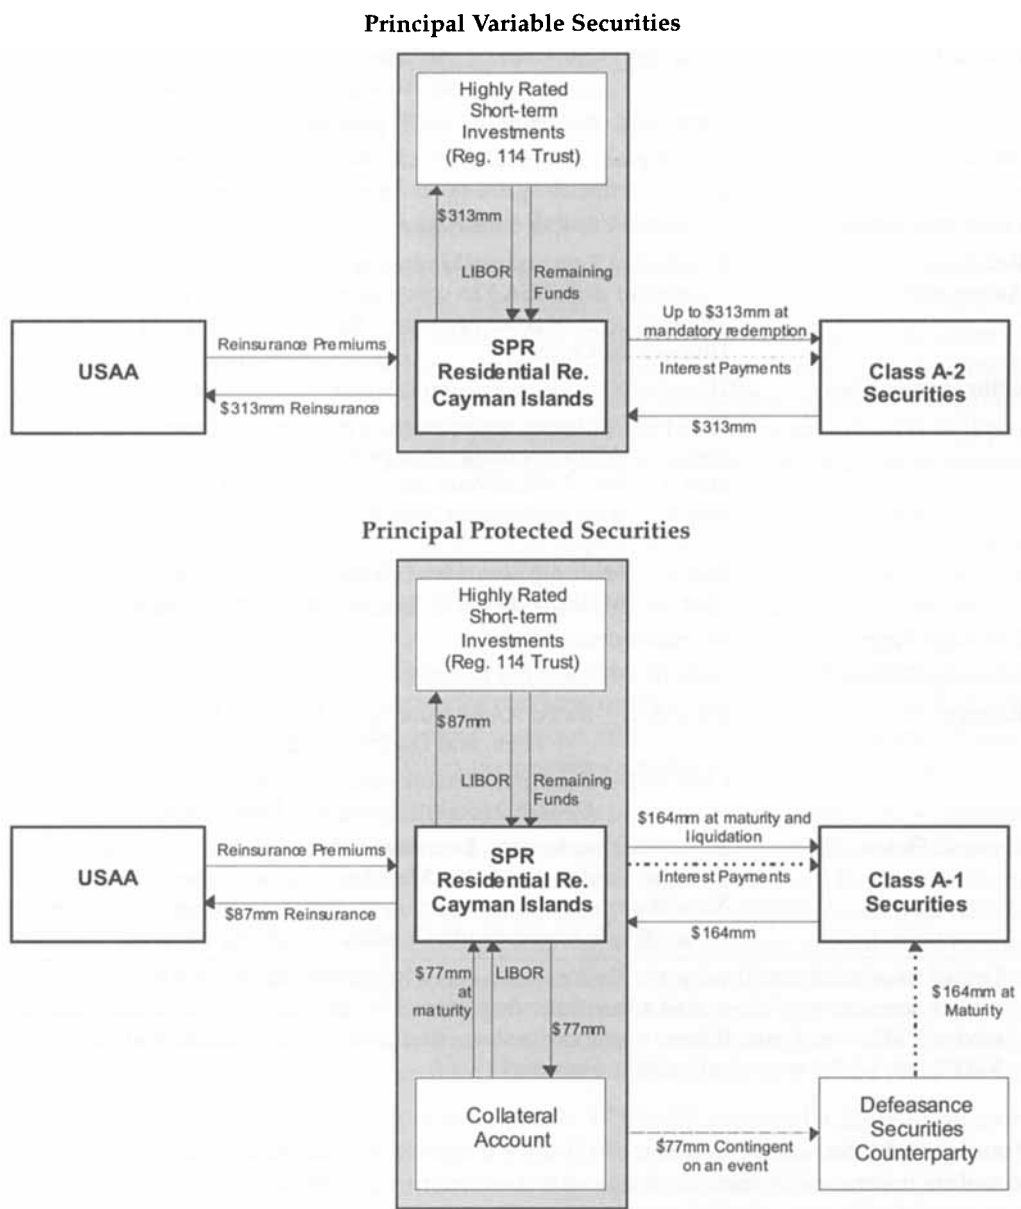
CAT event bonds are not the only vehicles for transferring risk. Swaps and forward contracts can be linked to CAT losses in much the same way as can bonds. However, these “derivative” contracts have an important shortcoming that bonds do not share—credit risk. Much like traditional reinsurance contracts, the cedent has a claim only on a

⁸ This section draws on Kenneth A. Froot and Mark Seasholes, “USAA: Catastrophe Risk Financing,” Harvard Business School, case no. N9-278-007, July 1997.

⁹ The reinsurance contract was 80 percent of the layer of losses in the range of \$500 million in excess of \$1 billion, deriving from a single hurricane. The structure of the Residential Re reinsurance contract is similar to that of standard reinsurance contracts. The main differences are that only one CAT-event loss can be claimed and that there are no reinstatement features.

¹⁰ To compare these percentages with actuarial losses, the risk of USAA sustaining losses in excess of \$1 billion is somewhat less than 100 basis points.

FIGURE 4
USAA CAT Bond Transaction Structure



Source: Kenneth A. Froot and Markus F. Mullarky, "Mid-Ocean Limited—Trading Catastrophe Index Options," Harvard Business School, case no. 9-298-073, April 1998.

TABLE 3
CAT Bond Contract Specifications

Obligor:	Residential Reinsurance Limited, a Cayman Island reinsurance company, whose sole purpose is to provide reinsurance for USAA
Amount:	Class A-1: \$164mm \$87mm principal variable \$77mm principal protected Class A-2: \$313mm 100% principal variable
Yield:	LIBOR plus: 575 basis points for the Class A-2 notes; 273 basis points for the Class A-1 notes. Interest is paid semiannually.
Loss Occurrence:	A Category 3, 4, or 5 hurricane
Reinsurance Agreement:	Residential Reinsurance Limited will enter into a reinsurance agreement with USAA to cover approximately 80 percent of the \$500mm layer of risk in excess of the first \$1,000mm of USAA's Ultimate Net Loss
Ultimate Net Loss:	Ultimate Net Loss = amount calculated in Step 6 (below) Step 1 All losses under existing policies and renewals Step 2 All losses under new policies Step 3 9% of the amount calculated in Step 1 Step 4 Add the amount from Step 1 with the lesser of Steps 2 & 3 Step 5 Multiply Step 4 by 1.02 for boat and marine policies Step 6 Multiply Step 5 by 1.02 to represent loss adjustments
Coverage Type:	Single occurrence'
Coverage Period:	June 16, 1997, to June 14, 1998
Ratings:	Class A-1: Rated AAAR/Aaa/AAA/AAA by S&P, Moody's, Fitch, and D&P, respectively Class A-2: Principal variable notes are rated BB/Ba/BB/BB by S&P, Moody's, Fitch, and D&P, respectively
Covered States:	Alabama, Connecticut, Delaware, District of Columbia, Florida, Georgia, Louisiana, Maine, Maryland, Massachusetts, Mississippi, New Hampshire, New Jersey, New York, North Carolina, Pennsylvania, Rhode Island, South Carolina, Texas, Vermont, and Virginia

* Unlike traditional reinsurance, the form of reinsurance offered by Residential Re was limited to one occurrence. If there was a hurricane that caused \$1,300mm of damage, the contract covered 80% of \$300mm. If there was another storm that produced USAA losses of more than \$1,000mm, USAA would no longer be covered.

counterparty, not a safe account. So it isn't surprising that securitized contracts that simulate insurance or reinsurance tend to be structured as bonds.¹¹

However, full collateralization is not the only way to reduce credit risk. Exchanges can reduce credit risk by marking to market frequently and by standing in between buyer

¹¹ In addition to credit risk issues, derivative contracts such as swaps and forwards run the risk of being classified as “insurance” contracts by insurance commissioners. “Insurance” contracts are highly regulated and are not considered appropriate for investors.

and seller to guarantee contract performance. With this problem solved, the exchanges are positioned to offer another benefit: the liquidity associated with trading standardized instruments. This is a benefit, indeed, because high-volume, low-cost trading in standardized, CAT event-linked contracts would seem an effective means of improving risk sharing.

To achieve standardization, exchanges have linked CAT event futures and options contracts to aggregated measures of CAT losses. The payoffs on these contracts are therefore tied to industry, rather than individual, losses. The Chicago Board of Trade is trading futures and options linked to the Property Claims Service (PCS) index, and the Bermuda Commodities Exchange is trading contracts linked to the Guy Carpenter Catastrophe Index (GCCCI).¹² Table 4 provides data on the exchange-traded contracts.

Standardized contracts have several important features worth noting. First, because of the aggregated nature of the indexes, they are not designed to indemnify—that is, to match the losses of any particular entity. As a result, any given insurer, reinsurer, or corporation ceding risk through these contracts will experience “basis risk” deviations between cedent losses and those paid by the contract. Sometimes when the cedent experiences a loss, the contract will not pay as much as the loss; sometimes it will pay more. All else being equal, this makes aggregate standardized contracts less attractive.

But the size of the basis risk varies considerably. It depends in part on the nature of the insurance portfolio being hedged and the choice of index. Clearly, lumpy portfolios that contain a few large industrial assets will not show losses that are tightly linked to total CAT losses across a large area like a state or region. Portfolios that contain homes evenly dispersed across an area will perform more like the area aggregate. Furthermore, most portfolio concentrations vary geographically. That is, market share changes considerably with location for insurance carriers. Indeed, market share varies not only across regions and states but also across counties and zip codes. Thus, the greater geographic disaggregation of the GCCCI index proves valuable for many portfolios.

Second, although standardized contracts generate basis risk for hedgers, they simplify life for investors. Imagine if investors were to buy diversified portfolios of CAT event risks by purchasing individual bonds. One at a time, investors would have to evaluate facilities such as Residential Re and the underlying risks of companies such as USAA. This is hard work and, even worse, redundant work. Each investor would need to analyze each individual risk separately. It is far more efficient for these risks to be bundled first into something that looks like an aggregated index before being marketed to investors. That way, each investor needs to understand only one kind of contract. The implication is that standardization makes it far less costly to distribute CAT event risks widely.

Along these lines, it is interesting to note that in 1998 USAA placed a layer of reinsurance through Residential Re in a new securitized transaction. This layer had similar risk characteristics to the one placed in 1997, also through Residential Re, yet the pricing of the 1998 bonds was considerably lower. On a comparable basis, the 1998 bonds provided a one-year return over LIBOR of 416 basis points, compared with 576 basis points for the 1997 bonds. This decline in return partly represents savings by investors and

¹² See the discussion below for more details on these indexes. [Note that the BCOE's trading was suspended in early August 1999 but might be resumed based on sufficient interest.]

TABLE 4

Comparison: The Proposed BCOE Catastrophe Linked Options and The Chicago Board of Trade PCS Options

	BCOE	CBOT
Membership		
Membership	<ul style="list-style-type: none"> • By application—no specific restrictions 	<ul style="list-style-type: none"> • By application—no specific restrictions
The Contracts		
Index used:	<ul style="list-style-type: none"> • GCCI Index 	<ul style="list-style-type: none"> • PCS Index
Reported metric:	<ul style="list-style-type: none"> • 1 unit = .01% industry loss-to-value ratio¹ 	<ul style="list-style-type: none"> • 1 unit = \$100,000,000 of industry losses
Geographic coverage:	<ul style="list-style-type: none"> • 7 geographic areas available: Northeast, Southeast, Gulf Area, Midwest, Florida, Texas (subject to data availability), and National 	<ul style="list-style-type: none"> • 9 geographic areas available: Northeast, Southeast, East Coast, Midwest, West, California, Florida, Texas, and National
Types of coverage:	<ul style="list-style-type: none"> • Aggregate losses • Single and second loss 	<ul style="list-style-type: none"> • Aggregate losses
Duration of loss periods:	<ul style="list-style-type: none"> • Semi-annual 	<ul style="list-style-type: none"> • Annual/Quarterly²
Payoff structure:	<ul style="list-style-type: none"> • Digital/Binary³ 	<ul style="list-style-type: none"> • Graduated depending on ending value of index
Operations		
Clearinghouse used:	<ul style="list-style-type: none"> • Bermuda Commodities Exchange Clearing House 	<ul style="list-style-type: none"> • Board of Trade Clearing Corporation
Clearing Member requirements:	<ul style="list-style-type: none"> • No specific restrictions 	<ul style="list-style-type: none"> • \$1 million of BOTCC stock • \$1.5 million of Memberships
Fees:	<ul style="list-style-type: none"> • \$20 round trip transaction fee per trade • LIBOR -100bp interest on cash margin (after first \$250,000) • 200bp/annum utilization fee on securities or letters of credit posted as margin 	<ul style="list-style-type: none"> • Roughly \$15 round trip transaction fee per trade⁴ • Short term U.S. Treasury interest on cash margin (less 50bp)⁴ • Zero utilization fee on securities or letters of credit posted as margin
Margin requirements:	<ul style="list-style-type: none"> • 100% posted up front and held through the term of the contract (\$5,000 per contract written) • Variation margin only posted if the market value of margin drops below \$5,000 	<ul style="list-style-type: none"> • Approximately 20% posted initially⁵ and adjusted daily on a mark-to-market basis • Clearing firms may require Members to post additional trading margin
Settlement	<ul style="list-style-type: none"> • Variable—up to 13 months from the end of the contract 	<ul style="list-style-type: none"> • Last business day of 12th month following end of loss period

¹ For example, if the industry is estimated to have suffered a loss to value ratio of 1.5 percent, the index will be published with a value of 150.

² California and Western contracts are annual; all others are quarterly. (National contracts available in both.)

³ Contract pays 100 percent of value (\$5,000) if the ending index value is above the option strike value, and pays \$0 otherwise.

⁴ Includes the brokerage fee charged by clearing firms to process Members' trades.

⁵ CBOT CAT option participants post margin equal to the greater of the normal margin maintained by other CBOT option traders (and based on a 20 percent volatility) and 20 percent of their maximum potential losses.

Source: BCOE, CBOT

intermediaries. They needed to expend far less time and energy understanding the transaction, having studied the same underlying package of risks one year earlier.

Third, it is worth distinguishing the role of standardization from that of the exchanges. Strictly speaking, standardization does not require an exchange, or for that matter even securities, to work. Traditional reinsurance contracts can (and are) standardized by adding triggers responsive to aggregated industry losses. However, the value of standardization is greater the more liquid the underlying contracts. With greater liquidity, standardized contracts can be sold and traded more widely. As with any other standard, the value is that everyone is on it.

THE FIVE KEY INGREDIENTS FOR A SUCCESSFUL CAT BOND ISSUE

When are CAT bonds and other securitized forms of risk financing likely to work best? Of the structures we have seen so far—contingent surplus notes, contingent equity puts, and CAT bonds—the bonds are the purest. They provide true risk transfer, which unlike pure financing alternatives does not merely prolong the day of reckoning. They avoid use of all-purpose equity financing, even on a contingent basis. They focus in on the CAT event risks at hand, promoting greater transparency and helping to reduce the costs of capital. They provide large capacity with zero-credit risk protection funded by competitive capital markets.

However, as we have noted, CAT bonds have weaknesses, and may not be appropriate in many circumstances. When, therefore, are CAT bonds likely to be economically advantageous over traditional reinsurance? What conditions need to be present in order for CAT bonds to be the preferred choice? In this section, we try to answer these questions by providing a list of ingredients—preconditions, really—that are likely to make CAT bonds the risk-transfer mechanism of choice.

Ingredient #1: The Retention Should Be Substantial

Like traditional reinsurance, CAT bonds need to cut possible losses into layers, so that each contract has an associated limit, or maximum possible loss. Low layers are those that are impacted frequently, as many CAT events generate at least a low loss level. High layers—like the USAA example above—are those that are impacted relatively infrequently, as few CAT event losses breach the contract retention (i.e., deductible). Most of the CAT bond transactions in Table 2 have low probability of loss and therefore translate into relatively high reinsurance layers. Indeed, the highest likelihood of loss among all the transactions is only 25 percent, and most have been completed at a probability of loss closer to one percent—a one-in-a-hundred likelihood of loss.

Several factors make this the case. First, there are institutional issues associated with fixed-income investing that bias choice toward higher layers. This is because many institutional fixed-income investors are restricted in their mandates: they must hold predominantly bonds that are rated; and they must hold bonds that have minimum credit ratings.¹³ The demand for higher-rated bonds is greater than for “junk” debt, which is deemed less than investment grade. In assigning a rating for bonds, credit-rating companies examine the probability of less than full repayment of principal and interest. A bond with an actuarial risk of loss of about two percent or greater is likely to generate a

¹³ Bond ratings are promulgated by firms such as Moody’s, Standard & Poor’s, Fitch, etc.

rating below investment grade, reducing the number and types of investors who can buy it.

Of course, with a little financial engineering, it is always possible to satisfy the rating agencies, improving the rating on relatively low-layer protection. The most common method is called “principal protection,” which was used for the USAA Class A-1 notes. Under principal protection, some of the money raised from investors is never put at risk. Instead, it is used to purchase long-term zero-coupon U.S. Treasuries at a deep discount. Since these securities will eventually pay their principal with certainty, CAT bonds can be designed to pay back their principal with certainty too. If there is a CAT event, investors lose money because they have to wait a relatively long time to get their principal back; and along the way, interest payments are at less than market rates. Using this feature, the Residential Re Class A-1 notes received the highest rating, AAA, while the Class A-2 notes were judged as being below investment grade at BB (see Table 3).


Naturally, there is a cost to this method: it waters down the CAT event risk embedded in the bond. Principal protection requires that much more than a dollar be raised for every dollar of risk transfer. The extra money finances the purchase of the principal protection. For example, in the USAA transaction, \$164 million of Class A-1 notes were issued. But of this amount, only \$77 million went toward backing USAA’s reinsurance. The rest went toward the purchase of zero-coupon U.S. Treasuries. As a result, the Class A-1 notes paid only 273 basis points above LIBOR, compared with 575 basis points for the Class A-2s. If USAA had wanted to protect a lower layer, this ratio would have been worse, with risk transfer making up an even smaller fraction of the money raised. Since marketing and distribution (and sometimes fee) costs of an issue depend critically on the amount raised, the costs of a low-layer CAT bond become problematic.

There is a second, entirely different reason why low layers are not naturally transferred through CAT bonds. Among insurers, reinsurers, and capital market investors, none like to bear the first dollar of loss from a risk. As a result, most contracts stipulate a deductible or retention. The reason is to give high-powered incentives to the original owner of the CAT risk. If a loss is declared, the owner pays the first dollars. This encourages good cedent behavior—better risk mitigation, loss-evaluation monitoring, and so forth. It also protects risk writers against private information cedents may have about their risk. The lower the layer, the worse the pool of cedents to come forward. Cedents who privately know the risk is bad are the ones who most want a low deductible.

These issues—formally referred to as moral hazard and adverse selection—are particularly important for CAT bonds because the capital markets are at arm’s length from the cedent. Traditional insurance and reinsurance methods bring the risk writer much closer. This is accomplished through ongoing information conveyed by brokers, by relationships, and by greater oversight and control over the cedent’s risk-assumption process.

Over time, as arm’s-length disclosure of information improves, lower-layer CAT bonds may become more common. But as we discuss in Ingredients 4 and 5 below, techniques for managing moral hazard and adverse selection are crucial, and disproportionately so for lower layers.

FIGURE 5
Comparison of Pricing and Riskiness of Different Types of CAT Risk Contracts



Source: Analysis provided by Risk Management Solutions.

The final reason that layers are likely to be high has to do with reinsurance pricing. So far, CAT bond returns have more or less mirrored those of the reinsurance market. This is shown in Figure 5, which compares the pricing of CAT bonds with traditional reinsurance contracts and options traded on the CBOT.¹⁴ And the pricing in the reinsurance market has been the most attractive in the highest layers. While this is apparent in Figure 5, it is clearly seen in Figure 6, which depicts the historical pricing of lower versus higher layers. Basically, large CAT risks are the risks most destructive to insurers and reinsurers because they create a large need for additional capital. The return for providing this additional capital is thus most generous relative to the risk. Thus, it is perhaps not surprising that CAT bonds have been targeted to compete with upper reinsurance layers.

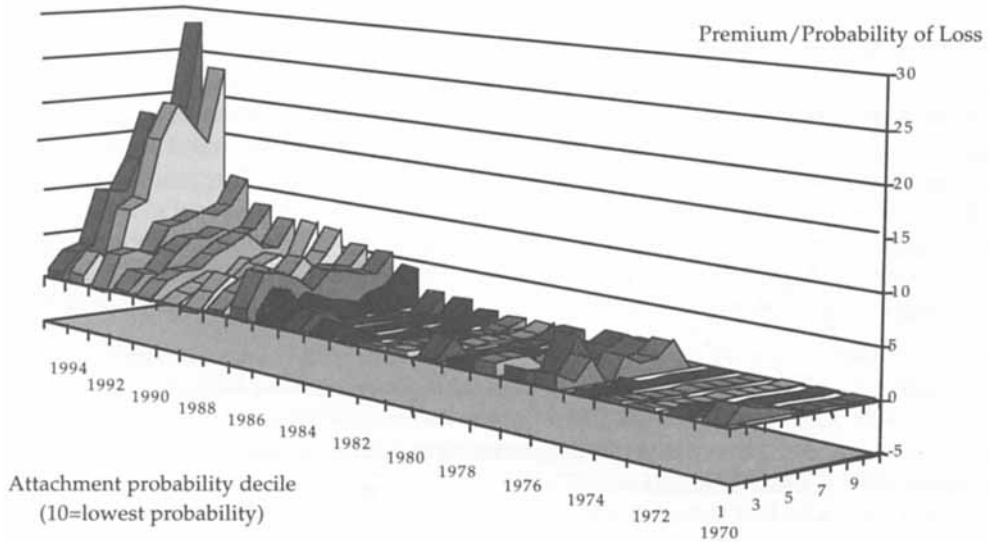
Ingredient #2: The Layer of Protection Shouldn't Be Too High

By itself, Ingredient #1 suggests that CAT bonds are relatively more attractive as the layer becomes higher. This seems natural, given the limited capacity of traditional markets. Why not go out into the vast liquid capital markets for funding not just the next higher layer, but three or four layers beyond that? The answer, of course, is as layers become higher, risk transfer ultimately becomes less economical. As the layer becomes too high, the probability of loss becomes too low. With very low probability of loss, the

¹⁴ The options which appear to have the lowest prices remain relatively illiquid, and the returns which appear on the chart are reduced by transaction costs. Transaction costs are not included in the returns for CAT bonds or reinsurance. Furthermore, the options are based on indexes of catastrophe losses, rather than net losses of individual insurers or reinsurers.

FIGURE 6

The Ratio of Premium to Actuarial Probability of Loss, by Layer



Source: Kenneth A. Froot, "The Limited Financing of Catastrophe Risk—An Overview," Harvard University, 1997.

premium ought to be very small. But a small premium will be insufficient to afford the costs of establishing the safe account, the SPR, or marketing and distributing the CAT bond. Moreover, the risks embedded in each bond, however small, are still unique. This forces investors to spend costly time understanding each risk. Very low risk CAT bonds simply do not pay.

To get a sense of how low premiums can go to be economical, it is worth considering several of the more precise costs of a CAT bond issue:

- **Underwriting fee.** Bankers will charge an underwriting fee, which could range between 50 and 250 basis points of the risk transfer amount, depending on the complexity of the structure and the quality of the underwriter. Given the state-specific regulatory nature of insurance, in some cases considerable travel expenses need to be paid by the cedent in addition to the fee.
- **Legal costs.** The legal costs of establishing the SPR and trust account are likely to be considerable. The legal work is complex, because the SPR must have the ability to write insurance or reinsurance protection and issue claims that are designated as securities and not as insurance or reinsurance policies. So it is not a standard (and less-expensive) type of special purpose vehicle used to create asset- or mortgaged-backed securities. For CAT bonds, the legal costs are likely to decline as the necessary legal expertise becomes more of a commodity. But currently this is not the case.
- **Cost of producing information on the risk.** Clients incur the cost of producing information on the risk for investors and rating agencies—in particular, fees to CAT

risk modelers. This cost can vary tremendously, depending on the risk. But it also can be a substantial expense.

- Fees to the rating agency. These are still somewhat higher than fees on normal corporate or asset-backed securities because of CAT bonds' still novel nature.
- Costs associated with capitalizing the SPV. These costs include directors' fees, licensing fees, claims and reserve consultant retainers, insurance fees, and so forth.

Taken together, these costs can easily exceed 150 basis points on a \$100 million issue. To see what this does to the economics of a high-layer CAT bond, consider a bond with a one percent risk of loss. The actuarially fair premium on such a bond would be 100 basis points. Even if investors were willing to buy bonds that paid only actuarially fair rates, the bare minimum charge—premium plus issue costs to the cedent—for such a bond would be 250 basis points. The cost is 2.5 times the minimum rate that investors are willing to receive.

It may be that cedent demand for protection is very great, but at some point the amount of protection gets too small relative to costs. At increasingly higher layers, the probabilities of impact fall. Clearly, at a layer so high as to have an actuarial risk of one basis point, the cedent costs would total 15 basis points—more than 150 times as large! Unless the demand for protection is extremely inelastic, there is too little cedent protection in such a contract to make it worth the premium.

Ingredient #3: The Dollar Amount of Risk Transfer Shouldn't Be Small

Because a large portion of the costs enumerated above are fixed, the 150 basis point threshold grows as the limit shrinks. A \$50 million issue would almost certainly require more than 200 basis points in cost. And the \$400 million Residential Re risk transfer was done for at least 100 basis points. These costs must either be paid in addition to the premium or must come out of investors' take. Clearly, small issues quickly become uneconomical.

The fixed-cost problem can be handled in part by using multiyear contracts. Although these are currently difficult to engineer, over time they will get easier. Today, participants are very unsure about how CAT risks and the associated premiums should evolve over time. Cedents may need to pay a higher premium to launch a multiyear issue successfully.

The fixed-cost problem can also be reduced through an increase in the market-wide volume of CAT bond issues. Legal, modeling, and investment banking fees are all prone to competition. A larger volume will support more competitors and allow fixed costs to be better amortized. Costs will therefore decline. But CAT bond volume will probably never grow large in comparison with that of other fixed-income markets such as corporate debt, asset-backed bonds, or mortgage-backed securities. So, at least for now, the prospect of razor-thin margins is a dim one.

Finally, over-the-counter transactions can be a partial solution to the costs above. The required level of underwriting fees, legal expenses, and modeling work can be reduced in small private transactions. This creates greater leeway in finding CAT bond possibilities. However, small private deals are by definition small and poorly advertised. The cost to a cheap distribution network can be that there is less competition among buyers.

Ingredient #4: The Loss Trigger Should Be Beyond Cedent Control

Cedent control of the loss trigger is known as moral hazard. Moral hazard tends to keep markets from developing. No one will come to trade when the outcome of a contract can be determined by the entity on the opposite side of the trade. In most insurance market instances, however, control is partial and often minor (e.g., workers compensation, product and environmental liability, errors and omissions). In such circumstances, it is not so much that moral hazard precludes exchange; it is more that the disadvantaged side must be paid sufficiently. The greater is the moral hazard—the control a cedent has of its losses—the greater the cedent payment must be.¹⁵

The first point to make about moral hazard is that CAT event risks provide relatively limited scope for it. CAT event risks are highly exogenous. So it is not surprising that they have been among the easiest insurance risks to trade. But, in spite of the exogenous nature of CAT events, cedents may nevertheless exert some degree of control over their losses. For example, the benefits of aggressive mitigation measures and loss cost and payment monitoring may not inure to the cedent once CAT risk has been transferred. As a result, the cedent may undertake less, increasing potential losses. Superior CAT bond structures are those that reduce opportunities for cedent control.

Dual triggers are a second way in which the scope for moral hazard is reduced. Many CAT bond triggers are linked to individual insurer losses (referred to ultimate net loss) as a first trigger, but there is almost always a second trigger as well.¹⁶ One obvious second trigger is that the claimed losses be associated with a CAT event as defined by Property Claims Service (PCS). In other instances, the second trigger adds a severity condition to the simple event trigger. Examples would be that a qualifying event must be a hurricane of class three or above, or an earthquake of magnitude 5.5, with an epicenter within a specific area.

A third way to reduce moral hazard is to eliminate the ultimate net loss triggers in a client trigger contract, leaving the company protection triggered off the size of industry losses associated with an event. A given company cannot manipulate industry CAT losses as much as it can its own CAT losses. So moral hazard with respect to industry losses is reduced.

Industry losses can be thought of as a kind of index to which individual company contracts can be benchmarked. The practical problem is that it is not easy to measure industry losses. Collecting data on losses that are directly related to a CAT event is difficult. It requires sophisticated information technology to classify claims (and paid losses) as specifically associated with a CAT event. Consolidating information across companies is difficult and slow, because each company collects and maintains information in its

¹⁵ The implications of moral hazard can be clearly seen in a recent CAT-bond transaction. The protection was intended to cover several lines of business and types of events. As a result, there was a concern that the cedent would have an incentive to grow those lines of business that were less attractive from investors' point of view. To limit the scope for moral hazard, the cedent agreed to a complicated set of restrictions—taking up some 80 pages in the offering memorandum. Moral hazard was reduced, but at the cost of considerable complexity and opacity.

¹⁶ When there is more than one trigger, contracts normally require both conditions are met before the cedent is entitled to obtain a recovery.

own way. As a result, existing indexes do not directly measure total industry-wide CAT event losses, even though the concept seems straightforward enough.

In practice, different indexes measure different things. An index published by PCS polls insurers in the aftermath of a CAT event, asking for estimates of their losses. This index has several virtues. Surveys can be conducted immediately after an event, making the result timely. The surveys are redone as claims come in, allowing for more accurate measures for a given CAT over time.

However, the PCS index has major disadvantages. Because it is based on surveys rather than hard data, it is difficult—indeed, impossible—to verify the results. Manipulation by reporting companies is possible, and error is likely. Moral hazard therefore remains a concern with the PCS index. In addition, large companies sustain a large proportion of losses. This means that estimation errors by large companies will remain important in the overall industry result. Finally, and most important, the index is computed on a state-by-state basis. States are large areas in comparison with CAT event footprints. Because many insurers' exposures are unevenly distributed throughout a state, individual insurer losses may correlate poorly with statewide losses. The index therefore provides less desirable protection for an insurer than would a contract linked directly to its ultimate net loss.¹⁷

A second approach to indexing is taken by the newly developed Guy Carpenter Catastrophe Index. It has several advantages as well. The GCCI is based on actual data from participating insurers. This removes guesswork and eliminates the scope for manipulation. These are advantages from the investors' perspective. The GCCI also makes it feasible to calculate disaggregated indexes on a zip code level. This is an advantage from the insurers' perspective, because it affords insurers better correlation with their own losses. Insurers can design a portfolio of zip code indexes to match their portfolio rather than rely on a one-size-fits-all statewide index.

The disadvantages of this index are a direct consequence of its use of hard data. Because the data collection requires time and considerable aggregation, GCCI results are available quarterly at pre-specified times. More frequent updates are not yet possible. Also, not all insurers report their data to the GCCI. Thus, the index uses a sampling approach, estimating loss-to-value ratios from participating companies within a zip code. These loss-to-value ratios can be aggregated across companies within any zip code, then grossed up to estimate losses within the zip code. Of course, since the sampling approach is what actually makes the hard-data index computation feasible, it may be better to think of it as an advantage. Indeed, statistical work suggests that the sampling errors are reasonably small, especially in view of the index's geographic disaggregation. Finally, the index currently reports only homeowner losses on atmospheric CAT events. Commercial/industrial losses and other events (such as an earthquake) are not covered.¹⁸

Some CAT bonds have been tied to index outcomes rather than ultimate net losses. For example, Swiss Re set up an SPR to provide it with \$112 million in California earth-

¹⁷ Contracts linked to the PCS index are traded on the Chicago Board of Trade.

¹⁸ Contracts linked to the GCCI index are traded on the Bermuda Commodities Exchange. See Kenneth A. Froot and Markus F. Mullarkey, "Mid Ocean Limited—Trading Catastrophe Index Options," Harvard Business School, case no. 9-298-073, April 1998, and The Guy Carpenter Catastrophe Index, January 1998, published by IndexCo (www.indexco.com).

quake protection. The protection and the bonds issued by the SPR were linked to industry losses—measured by the California PCS index—associated with a single earthquake in California. For more details, see Table 1.

As already mentioned, indexes help reduce moral hazard on the part of the insurer. They also help standardize CAT event risks, helping investors avoid analyzing the properties of each company's individual CAT event risk. These features should, in principle, make index-linked protection cheaper (for the cedent) than indemnification through an ultimate net loss trigger. Currently, because there have been so few transactions, there is not much evidence that indexes are cheaper. But as investors begin to use indexes as benchmarks, a cost disparity should emerge.

Ingredient #5: The Loss Trigger Should Be Symmetrically Transparent

The last important ingredient concerns asymmetric information. It is clearly difficult to cede a risk to another party when it is known that the cedent has superior information. In that case, the buyer (the assumer) is a sheep and the seller a wolf. Cedents who wish to lower the costs of protection have an interest in credibly providing information and transparency about the risk to be sold. However, investors will be concerned that some cedents will not worry so much about the cost of protection. They worry that those cedents who view the cost as cheap are those who know that their risk is worse than appreciated. The effect becomes worse as the cost of protection is higher. Thus, the greater the information asymmetry, the greater the "adverse selection" of transactions against the investor.

There are several points to be made about adverse selection. First, indexes can help reduce informational asymmetries in much the same way that they reduce moral hazard. Investors may credibly know as much about index losses as any particular company would know. A company with risks that are worse than appreciated by investors has no particular incentive to link its protection to an index. The reason is that if the company loses more than the index, the company would suffer, not investors.

The second point about adverse selection is that it depends on an asymmetry in buyer-versus-seller information and not on the level of available information. There are no sheep without wolves. For example, it may be impossible for anyone to know how the courts will handle certain environmental or health liabilities in the future. But as long as cedents know as little about the risks as do investors, protection can be obtained, even if very little is known in an absolute sense about the risk. Indeed, earthquakes and hurricanes are this way. We know very little about when and where the next big event will occur. Yet because much of the ignorance is common to buyers and sellers, protection for such events can be obtained.

The third point is that adverse selection can be mitigated in several ways. First, a high level of disclosure is helpful. Given a cedent's information, it is on average better to raise buyers' information levels. Second, companies can incur more damage by being seen as preying on sheep than by paying the cost of protection. For example, faulty products ruin a company's reputation and can severely reduce its market value. Managers, whose compensation is increasingly tied to stock price performance, have little incentive to produce faulty products, because writers of product liability insurance do not bear the full cost of damage to a firm's reputation. They would prefer to buy product liability insurance for unforeseen matters, but to do everything possible to mitigate the scope for, and effect of, product imperfections.

LOOKING TO THE FUTURE

Several of the five key ingredients for a successful bond issue are likely to evolve over time. For example, the importance of having a high retention (Ingredient #1) will likely decline. The cause will be a reduction in institutional barriers to issuing lower-rated CAT bonds. There is, after all, nothing that says low-rated CAT bonds must pay less than high-rated corporate bonds. Certainly, low-rated CAT bonds will offer greater yields and commensurately greater risk. But, as with any CAT risk placed into a fund, much of the risk is diversifiable. This is true even within catastrophes: combining the risks of California earthquake, East Coast hurricane, and North American freeze together with those from outside the U.S. can reduce pooled risk enormously. Funds whose mandates allow them to buy low-rated bonds—perhaps dedicated entirely to CAT—will create a market for securitization at lower retention levels.

This institutional evolution is likely to be important in a number of ways. The costs of issuing CAT bonds are in large part a function of the volume of issue. If there is little volume, few investment banking groups specializing in these bonds can be supported, and competition will be low. Few investment managers will specialize in the area, so the general level of investor expertise will be low. This will make funds generally less willing to purchase new risks and will raise the cost to potential issuers considering the capital markets. For these reasons, the market for CAT bonds is subject to increasing returns—it will run smoother and more efficiently if the volume of CAT issues is relatively high. And, of course, this is self-reinforcing—if the volume of CAT issues is higher and costs are lower, more companies will want to use them.

We are already seeing additional investment pools today that will specialize in taking insurance risk. A first group of these will look like traditional reinsurers—raising equity from investors, forming a company, and using the money to write insurance risks. A second group will do something similar, but specifically from tax havens such as Bermuda. This permits an investment that allows funds to compound tax-free, and tax-free compounding is valuable to investors. These firms will pursue investment strategies that are particularly tax-disadvantaged, in order to maximize the value of the tax-free compounding they can offer. Note that because such firms can add value by virtue of tax-free compounding of their investments, they do not need to add much—or even any—value in their insurance/reinsurance operation.¹⁹ Thus, these firms will be a potent source of competition, helping drive down prices of insurance and reinsurance. This is one way of increasing the amount of low-cost capital willing to provide CAT event risk capacity.

A third group will emerge looking more like mutual funds, where investors put money into a fund (not a company) and own shares of the fund assets. For open-ended funds, these shares will be priced daily on a marked-to-market basis. The fund assets go toward buying CAT bonds and other insurance risk securities. Such fund structures will help to reduce the cost of capital. However, this structure also reduces efficiency somewhat, since each fund investment will be on a fully collateralized basis, and the “leverage” employed by insurers and reinsurers is lost. Still, the fund structure will be very important for providing a potentially large and rapidly expandable pool to discipline prices.

¹⁹ Under U.S. law, it is the existence of the insurance/reinsurance operations that allows the assets to qualify for tax-advantaged status.

Together, these three institutional prototypes will provide easier access to capacity when needed. By historical standards, prices will be low. And, in spite of any uncertainty about the absolute level of issuance, a more basic truth emerges: because of these developments, the insurance cycle is fundamentally and permanently changed. In the past, large events depleted balance sheets and triggered enormous increases in prices. Afterward, the high prices encouraged an increase in the amount of capital willing to write risk; but the increase occurred slowly, over several years. As capacity increased, prices fell again.

These new sources of insurance risk transfer undercut this cycle by disciplining prices. Capital will rush in quickly if prices rise substantially, and this keeps prices from rising in the first place. Prices may rise after a CAT, but not nearly to the extent that we have seen in the past. The days of large cycles and extremely profitable underwriting opportunities appear to be coming to a close. Competition for capacity is coming to insurance and reinsurance, just as it has come to every other financial services business over time.

If insurers and reinsurers are to succeed in the world just described, they must reduce their costs of capital. One way to do this is to make their businesses and balance sheets more transparent. Increasingly, investment portfolios of these firms are marked-to-market. Bad news is reflected immediately. This means that problems cannot as easily be swept under the rug and be allowed to build, unobserved, for a long period of time (e.g., the banking crisis of U.S. savings and loans, Japan's real estate loans, etc.). By extension, marking-to-market of liabilities will help to shrink surprises, reducing moral hazard and adverse selection. Insurers and reinsurers will increasingly be under competitive pressure to do this. Interim trading of CAT event risks will facilitate—even force—this development.

Thus, CAT bonds are going to be a powerful force even if we don't see many dedicated bonds issued. The analogy is the way Saudi oil output affects U.S. pump prices, even though the United States buys mostly from Venezuela and Mexico. In fact, as CAT and other insurance risks become more transparent and familiar to more investors, there will be little need for dedicated CAT event bonds. Corporations will be able to add CAT event clauses to their standard debt issues and get competitive pricing on the combined credit and CAT risks. This kind of development will further reduce the costs of issuing CAT-linked notes while increasing the amount of risk transfer.

The basic conclusion here is that the vision of a world that switches wholesale from insurance to CAT bonds is a *reductio ad absurdum*. CAT risk will be transferred more efficiently as we go forward. The insurance cycle will, for all intents and purposes, be gone. But the transfer will occur in many ways—some obvious, some creative. And the CAT bond will be only one of many ways to skin the CAT.

The following language was included with this article when originally prepared as a special report by Marsh & McLennan Securities Corp.:

The information contained herein is based on sources we believe reliable, but we do not guarantee its accuracy. Past performance of financial instruments, indices, or events cannot guarantee future outcome. This document is not an offer to sell, nor a solicitation of an offer to buy, any financial instrument or insurance or reinsurance program. The material, data, and/or methodology used in this special report are proprietary to Marsh & McLennan Securities Corp. and Guy Carpenter & Company, Inc. The dissemination or use of this report without Marsh & McLennan Securities Corp. or Guy Carpenter & Company, Inc.'s express written permission is prohibited.