Analyzing Insurance-Linked Securities

EDUARDO CANABARRO, MARKUS FINKEMEIER, RICHARD R. ANDERSON, AND FOUAD BENDIMERAD

EDUARDO CANABARRO is a vice president at Goldman, Sachs & Co., in New York.

MARKUS FINKEMEIER is a vice president at Goldman, Sachs & Co., in New York.

RICHARD R. ANDERSON is the chief actuary at Risk Management Solutions in Menlo Park, California.

FOUAD BENDIMERAD is technical director at Risk Management Solutions in Menlo Park, California.

Insurance companies are in the business of assuming risks from individuals or companies. They manage those risks by diversifying over a large number of policies, perils, and geographic regions. A particularly difficult problem is the management of risk from high-severity, low-probability events (*catastrophe risk*, or *"CAT" risk*), such as that posed by major earthquakes or hurricanes. The risk from low-severity, high-probability events (for example, auto collision coverage or medical insurance coverage) can be diversified by writing a large number of similar policies.

Suppose the insurer charges a premium equal to the expected average annual loss and has a very large number of policies. By the law of large numbers, it can expect to pay out approximately this amount in claims each year. It can then earn a stable profit on the investment income from the premiums. However, the law of large numbers breaks down when confronted by infrequent and severe catastrophe risks — losses can occur across an entire pool of policies simultaneously. Events that occur with annual probabilities of the order of 1% or less — i.e., about once every 100 years or more — are sometimes called *super-CAT events*.

An important tool to manage catastrophe risk is reinsurance. Reinsurance is the means by which insurance companies transfer their own portfolio risk to other reinsurance companies. Exhibit 1 displays the structure of an *excess of loss* (*XOL*) reinsurance contract. Against the payment of a premium, the reinsurer agrees to provide the *ceding insurer* (*cedant*) with protection against a *layer of losses* above a certain level.

For example, assume the reinsurer provides "$50 million of protection in excess of $100 million, with 10% coparticipation." This means that the reinsurer makes payments to the insurer if its losses are above the *attachment point* of $100 million. If the losses are between $100 million and the *exhaustion point* of $150 million, the reinsurer pays 90% of the amount exceeding $100 million. If the loss is larger than exhaustion, the insurer receives the maximum amount of 90% × $50 million = $45 million. *Proportional reinsurance*, or *quota-share*, is an alternative to XOL contracts in which a reinsurer provides capital on a proportional sharing basis. In this type of contract, the cedant is reimbursed for a fixed percentage of its losses in return for ceding a fixed percentage of premiums.

Reinsurance companies manage their risk by an even wider and more global diversification. The reinsurance market is large and well established, but the total capital and surplus of reinsurers is dwarfed by that of insurers. In the late 1980s and early 1990s, several large losses from U.S. catastrophes (Hurricanes Hugo and Andrew and the Northridge earthquake) put large strains on the capacities of the reinsurance markets (estimated to total $50 billion-100 billion).

Although the situation has eased in the

last few years, prices in the reinsurance market remain very volatile and could potentially increase again, perhaps significantly, after one or more major future catastrophes. Innovative mechanisms have been developed in recent years to stabilize pricing and coverage by transferring risk, in security or derivative form, into the much larger pools of investment capital available in the global capital markets. Since catastrophe risk is inherently uncorrelated with financial markets, such instruments are potentially attractive to many types of investors. Thus, an integration of the reinsurance and the capital markets can be mutually beneficial for both insurers and investors, providing insurers with reinsurance coverage at reasonable and stable prices, and offering attractive securities to investors.

STRUCTURE OF INSURANCE-LINKED SECURITIES

We now examine in detail how the transfer of risk to the capital markets can be achieved using insurance-linked securities (sometimes referred to as catastrophe bonds, or *CAT bonds*). Exhibit 2 illustrates the typical structure of a catastrophe bond. A *special-purpose vehicle* (*SPV*) is created, which performs two functions: It provides reinsurance to a ceding insurance (or reinsurance) company, and it issues securities to investors.[1] At the beginning of the risk period, the SPV collects funds from the investors, which are deposited in a trust. At the same time, the cedant pays a premium to the SPV in exchange for the reinsurance that the SPV underwrites. This premium and the investment income on the trust provide the funding for the interest payments the investors receive (an investment return, generally LIBOR plus a spread). If no *trigger event* causing losses above the attachment point of the reinsurance contract occurs during the risk period, investors receive back their principal at maturity. However, if a loss above the attachment point of the reinsurance contract does occur, the funds are first used to make payments to the ceding company under the agreement, leading to a partial or total loss of principal to investors.

EXHIBIT 2
Structure of a Catastrophe Bond

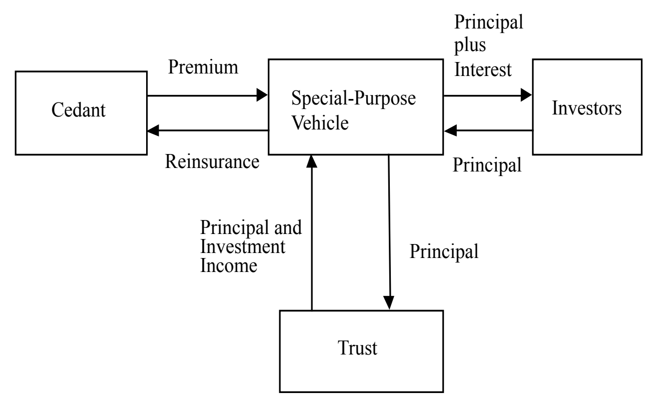

EXHIBIT 1
Structure of an Excess of Loss Agreement

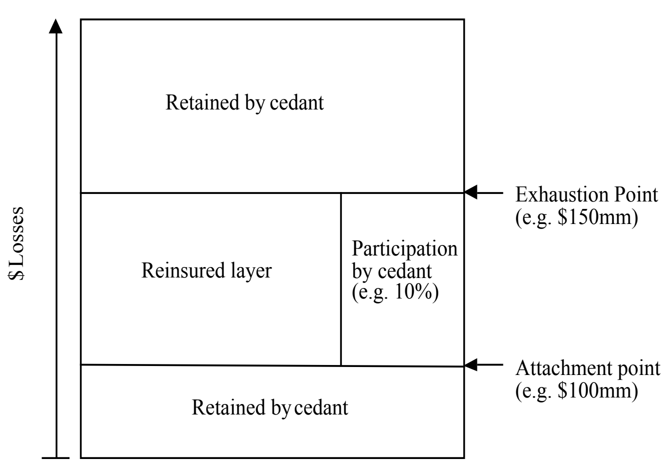

An important issue is the definition of losses in the contract between the cedant and the SPV. If the contract is based on the actual loss experience of the cedant's own book of business, this is known as an *indemnity-based* contract, and closely resembles a traditional reinsurance program. Alternatively, loss payments can be based on an *index* (for example, of industry loss experience). An indexed structure can make it easier for investors to analyze the risk, because they no longer need to understand the details of the cedant's business.

However, the cedant can be exposed to *basis risk*, to the extent that its own exposure differs from that of the index used to determine the payoff of the contract. It is possible that, in a given event, the cedant experiences large losses, while the losses in the index — and thus the payments to the cedant — are relatively small. Because of this basis risk, the cedant may want to pay less of a premium for such coverage.

A final possibility for the contract structure (and mirroring security) is to use the physical parameters of the natural hazard, such as the magnitude and location of an earthquake, as the trigger (*parametric structure*). In the traditional reinsurance market, indemnity-based contracts predominate, but indexed structures, commonly referred to as *industry loss warranties (ILWs)*, are sometimes used to alleviate concerns about underwriting quality if there is limited capacity, or to eliminate the necessity of disclosure by the cedant.

Insurance-linked securities have often been offered in *"principal-at-risk"* and *"principal-protected"* tranches. In the case of a principal-at-risk bond, the full principal amount is available to support the contract provided to the cedant in case of a catastrophe, and thus the full principal amount can be lost. The coupon is typically guaranteed, such that there is a minimum "recovery value." For any such security, it is possible to structure a principal-protected tranche, where some portion (typically about 50%) of the principal is available as risk capital of the SPV. The other 50% are held in a separate, protected account so that these funds can be used to ensure that the full principal can be repaid.

One way to achieve this protection of invested principal is to enter into a derivative contract at a small cost, which guarantees that in the event of adverse loss experience, the protected 50% will be sufficient to purchase zero-coupon Treasury bonds that will accrete to the full face value of the principal within a predetermined period of time (e.g., ten years).

This principal-protected tranche has been structured and offered to enable bond buyers who have restrictions on the amount of non-investment-grade securities they can purchase to participate in the new asset class. While principal-at-risk CAT bonds have typically been rated B or BB, the principal-protected bonds carry up to a triple-A rating with respect to repayment of principal.

In Exhibit 3, we give an overview of the major insurance-linked security transactions and their respective structures and offering spreads.

CATASTROPHE MODELING

Assessment of Natural Catastrophe Risk

For determining a fair price, and for investors to be able to judge how insurance-linked securities fit into a portfolio, it is crucial to have a good understanding of the risk of potential losses. Both the expected loss and the probability distribution of potential losses are important in this respect.

In Exhibit 4, we sketch a continuum of possible approaches to this problem, using hurricane risk as an example. One approach is to construct probability distributions for future insurance losses based on records of historical insurance losses. Such records are available, for example, from Property Claims Service (PCS) back to 1949. These historical losses can be adjusted for the effects of inflation, and we can estimate adjustments for population growth and movement.[2]

The use of historical loss data for the estimation of future losses, however, is very problematic. Too many variables affecting the size of these insurance losses have changed materially. Population growth has been disproportionately high in coastal and seismic areas within hurricane- and earthquake-prone states such as Florida and California. Building materials, construction designs, and property values have changed. Factors mitigating CAT losses, including building codes and their enforcement, and warning and damage management strategies, have also changed, partially in response to recent catastrophes. Furthermore, underlying insurance policies may have changed materially, with respect to both the number of insured properties and the details of the coverage, such as maximum insured loss, coinsurance, and deductibles.

To overcome these problems, we can instead estimate the effects that historical storms would have if they were to occur today, given today's buildings and insurance policies. This leaves us with various problems related to the limited size of the sample of historical events ("tail problem"). This problem is even more serious for earthquakes. Large earthquakes in a given location may have recurrence periods (i.e., estimates of time between events) of several hundred years. Hurricanes are more frequent, and thus we have a larger database of historical events. Between 1899 and 1997, almost 160 hurricanes made landfall in the continental U.S.,[3] or an average of about 1.6 hurricanes per year, though large or very severe hurricanes are rare in any given geographic location.

The small-sample problem can be attenuated by making parametric assumptions about the probability distribution of the characteristics of the hurricane. We can assume, for example, that the number of hurricanes at a given location and the parameters of each storm (such as the forward velocity or the central barometric pressure) are independent. Then we can use the set of characteristics of all historical storms to build the probability distributions for each of these parameters, and proceed to build synthetic hurricanes by sampling from these distributions independently.

This approach, which combines the use of historical data with certain parametric assumptions, is the main

EXHIBIT 3

Major CAT Bond Transactions*

Issue Date	Special-Purpose Vehicle	Cedant	Total Size ($mm)	Risk Capital ($mm)	Risk Period	Risk	Structure	Spreads (L = LIBOR) (bp) for Different Tranches
12/96		Hannover Re	100	100	5y	Portfolio	Indemnity/ Loss Ratio Indexed	(Dividends)
12/96	George Town Re	St. Paul Re	68.5	45	3y/10y	Portfolio	Indemnity/ Proportional Participation	(Dividends)
06/97	Residential Re '97	USAA	477	400	1y	Single East and Gulf Coast Hurr.	Indemnity	L+582/L+276
07/97	SR Earthquake Fund	Swiss Re	137	122	2y	CA Earthquake	Index	2y UST + 625, +475/+280 or L+259
11/97	Parametric Re	Tokio Marine	100	90	10y	Tokyo Area Earthquake	Parametric	L+436/L+209
03/98	Trinity Re '98	Centre Solutions	83.6	72	10m	Florida Hurricane	Indemnity	L+367/L+ 157
06/98	HF Re	Continental Casualty	90	90	6m	NE Hurricane	Index	T-bills + 375
06/98	Residential Re '98	USAA	450	450	50 wks.	Single East and Gulf Coast Hurr.	Indemnity	L + 404
06/98	Pacific Re	Yasuda Fire and Marine	80	80	5y-7y	Japan Typhoon	Indemnity	L+375 for a high layer protection, switching to L+963 for a lower layer following a drop-down trigger event
07/98	Mosaic Re I	F&G Re (St. Paul Re)	54	45	50 wks.	Various U.S. Perils (largely Florida hurricane)	Indemnity	Class A Notes (senior principal-at-risk): L+440 Class B Notes (subordinated principal -at-risk): L+820 units (principal protected): L+215
08/98		XL Mid Ocean Re	100	100	1y	U.S. Hurricane and Earthquake	Indemnity	L+418 / L+598
01/99	Gemini Re	Allianz	150	150	3y-6y	German Windstorm	Indemnity	49 annual option premium. After the first event, issuer can put bonds with L+833
01/99	Trinity Re '99	Centre Solutions	56.615	54	1y	Florida Hurricane	Indemnity	L+423 / L+177
01/99	Mosaic Re II	F&G Re	45.7	45	1y	Various U.S. Perils, with Very Little Florida Hurricane	Indemnity	Class A Notes (senior principal-at-risk): L+406 Class B Notes (subordinated principal-at-risk): L+836 units (principal protected): L+193
03/99	Domestic, Inc.	Kemper	100	100	3y	New Madrid Earthquake	Indemnity	Equity tranche (subordinated): dividends. Debt tranche (senior): L+374
05/99	Concentric Ltd.	Oriental Land	100	100	5y	Tokyo Area Earthquake	Parametric	L+314
05/99	Circle Maihama Ltd.	Oriental Land	100	100	5y-8y	Extension and Credit Risk	Earthquake-Contingent Financing	L+76
06/99	Residential Re '99	USAA	200	200	1y	Single East and Gulf Coast Hurr.	Indemnity	L+372
06/99	Juno Re	Gerling	80	80	3y	SE U.S. Hurr.	Indemnity	L+426

*This table includes transactions since December 1996 exceeding $40 million of risk capital that was broadly distributed. Note that we have multiplied spreads over LIBOR by #days/360, and therefore the numbers displayed differ from the spreads quoted in the offering circulars. For transactions with risk periods of more than one year, #days = 365, the number of days *per year*. For transactions with a risk period of less than one year, #days = the total number of days in the risk period. For a detailed explanation, see the appendix.

EXHIBIT 4

A Continuum of Possible Approaches to Hurricane Risk Assessment

Approach	Issues/Features
Use time series of historical losses or loss ratios (e.g., PCS data).	Because of inflation and population growth, future losses can be expected to be much larger than past ones.
Correct historical losses for inflation and population growth.	Population growth has been disproportionately large in high-risk areas. Buildings, insurance, and mitigants have changed materially.
Simulate the effects of historical storms on current exposures.	Very small data samples in smaller geographical areas and for severe events (tail problem).
Estimate the correlation between the main hurricane parameters, and create synthetic storms by sampling from (smoothed) historical distributions for each of these parameters.	Requires selection of specific types of storms and weightings to include in the distributions. The correlation estimates might not be correct.
Create a physical model of hurricane formation and progression, simulating the interaction of the atmosphere, ocean waters, and the surface over land.	The system ocean-atmosphere-land is extremely complex, not very well understood, and thus very difficult to model.

approach used by catastrophe modeling firms such as Risk Management Solutions (RMS), Applied Insurance Research (AIR), EQECAT, Risk Engineering, and others.

At the extreme, one could create a computer model of the physical process of the formation and progression of hurricanes (or earthquakes). This model would have to represent the complicated interaction between the ocean, the atmosphere, and the land surface after landfall. Any accurate representation of this complex system would be extremely difficult.

Stochastic Approach to Risk Assessment

Exhibit 5 displays the structure of the computer modeling approach to catastrophe risk assessment, which consists of three modules or steps. The first module, the *hazard module*, defines the range of potential stochastic events representing the natural hazards (such as earthquakes and hurricanes) in terms of their physical characteristics and their probabilities of occurrence. The inputs of this module consist of historic data (such as storm catalogues or earthquake history), site characteristics (such as terrain roughness or soil conditions), scientific knowledge, and expert opinion. The outputs of the hazard module consist of a set of stochastic events and their determinant characteristics at each location, such as wind speeds for hurricanes or ground motion for earthquakes.

The second module, the *damage module*, then determines the damage caused by a specific storm or earthquake to the structures, such as houses, industrial installations, or infrastructure. One way to arrive at functions that map the natural hazard determinants into damages (*damage functions*) is to apply structural engineering expertise. Another, potentially more reliable, method used by modeling firms employs large databases of actual claims experienced by insurers in various catastrophes to construct damage functions by regressing the claims against the historical natural hazards. The output of the damage module is the estimated damage caused by that specific event to each of the various structures.

The third module, the *financial module*, then applies these damages against the insurance or reinsurance contract specifications to determine the impact of the estimated event (if any) on the relevant policies and calculate the ultimate financial losses.

Hurricanes, Typhoons, and Tropical Cyclones

Tropical cyclones are a class of intense cyclonic wind systems — i.e., wind systems that rotate in a closed weather circulation — that develop over tropical waters. They are called "hurricanes" if they occur in the West Atlantic, "typhoons" in the West Pacific, and generically "cyclones" in the Indian Ocean. They consist of a calm core (the eye of the storm) surrounded by high walls of thick clouds and rotating winds. In addition to intense

EXHIBIT 5

Illustrative Framework for Catastrophe Modeling

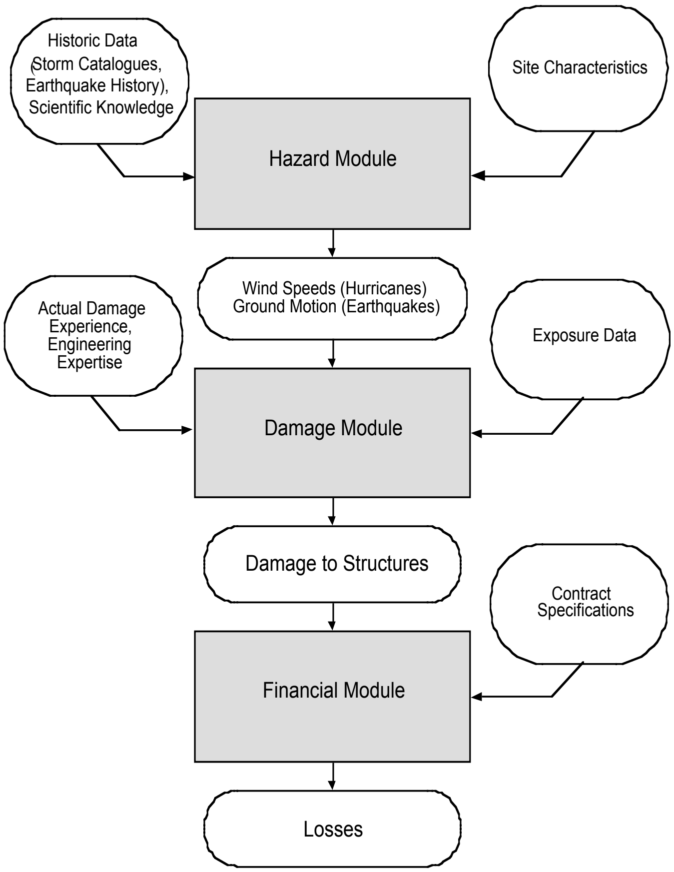

winds, they can create very heavy precipitation and tornadoes. In the Northern Hemisphere, they usually occur during the months of June through October.

When hurricanes make landfall, they can cause loss of life, extensive damage to buildings and other infrastructure, crop and forest destruction, and water contamination. Storm surge refers to the rising of the level of the ocean water along coastlines associated with hurricanes that can cause widespread flooding.

Hurricane intensity is related to the difference between the low pressure within its eye and the ambient pressure far away from it. Thus its intensity can be classified either in terms of the central barometric pressure (the lower this is, the more intense the hurricane is) or in terms of the wind speed. The Saffir-Simpson scale divides hurricanes into five categories, ranging from the moderate Category 1 to the rare and catastrophic Category 5 hurricane.

Earthquakes and Earthquake Modeling

Earthquakes can be generated by a sudden dislocation of large rock masses along "fault lines" — fractures within the crust of the earth. Geologists believe that the crust of the earth is composed of several crustal plates, and the relative movements with respect to one another cause earthquakes. Earthquakes can occur "interplate" — i.e., along their edges — where they may collide, slide past one another, or pull apart from one another. In "intraplate" earthquakes, forces are transmitted from the edges of the crustal plates and result in quakes in their interiors.

The main parameters of an earthquake are its location, depth, fault rupture plane, and magnitude. The location, or "epicenter," of an earthquake is defined as the point on the surface of the earth directly above the "hypocenter" — the initial point of rupture within the earth. The "depth" is the distance between the hypocenter and the epicenter. A quake starts at a single point (the hypocenter), and then propagates through a fault rupture plane. The area of this fault rupture plane is an important determinant of the magnitude of the earthquake, and the distance of a location to the plane of rupture is the relevant distance in the damage analysis.

The magnitude of an earthquake is a measure of the total energy that is radiated from its rupture plane. Various common magnitude scales are in use, such as the moment magnitude, the surface wave magnitude, the body wave magnitude, and the local magnitude scales. These scales can empirically be related to one another. The media often report earthquakes using the "open-ended Richter scale." This scale was developed for a specific type of seismograph that is no longer in use. Richter magnitudes are typically local magnitudes, but do not imply a specific scale.

The concept of "recurrence interval" of earthquakes, i.e., the characteristic period of time between events of similar magnitude, has to be distinguished from the concept of "return period" of hurricanes. For hurricanes, probabilities are generally assumed constant from year to year (although some forecasters do make them conditional, for example, on El Niño). Thus, the return period is simply the inverse of the annual probability, which is time-independent. For earthquakes, however, scientists generally believe that at least certain faults have time-dependent probabilities. They believe that such faults build up stress at a constant rate and will release it periodically when it reaches certain levels. This means that the probability of an event occurring at a given location increases as the time since the last event at this location increases.

Recurrence intervals increase with the magnitude of the earthquake and, for very large earthquakes, can be in the range of hundreds or even thousands of years. To give two examples, the recurrence interval for the 1906 San Andreas earthquake (magnitude 7.9) is estimated to be about 210 years, and that of the 1992 Landers quake (magnitude 7.3) to be about 5,000 years. Available historical data for such large earthquakes are very limited and cannot be relied upon as the only source for estimating future probabilities. Geological and geodetic information can significantly enhance the estimation of recurrence intervals of large events. This includes data on the slip rates of faults, and trenching, which can provide information on earthquakes that occurred in prehistoric times ("paleoseismology").

Apart from hurricanes and earthquakes, models have also been developed for other perils such as tornadoes, straight-line windstorms, European windstorms, floods, and hail.

Issues in Assessing the Reliability of Catastrophe Modeling

Issues related to differences between the various risk models and the reliability of their respective risk estimates include:

- It is important to understand which risk components have been included in the model estimate and which have not, and if they agree with the scope of the relevant insurance contracts. This relates in particular to *fires following earthquakes* and to *storm surges and floods after hurricanes*.
- Two important effects that are typical optional components within a model are *demand surge* and *loss adjustment expenses*. Demand surge describes the potential of temporary inflation after a major catastrophe resulting from localized increased demand for building materials, repair workers, etc. Loss adjustment expenses describe legal and operational expenses incurred by the insurer in connection with the settlement of claims. Both effects are typically expressed in the form of factors applied to the original loss estimate.
- The *quality of the exposure data* provided by the insurer is very important. The main issues are accuracy, geographical resolution, and the availability of detailed information on the insured structures and on the individual insurance policies. First, the data might contain outright errors, such as incorrect address, building type, or policy details. Second, data provided with a street address and a zip code, and with clearly defined individual policy structures, result in a better estimate than data aggregated to a county or state level. Third, details about the insured structures such as building type and year of construction may or may not be available. Fourth, details about the individual policies such as deductibles, coverage amount, and exclusions are not always available. The resolution and quality of the exposure data have a direct impact on the accuracy of the risk analysis, as either actual information or assumptions must be provided to the model.
- Hurricane modeling generally assumes *a global climatic system that was stationary* over the last 100 years. If, however, ocean water temperatures are slowly increasing, this assumption may not be accurate. The current methodology will not reflect the impact of such a potential warming of ocean waters.
- *Conditional analyses* — modeling estimates that consider potential differences between specific years (such as El Niño and La Niña years) — are not typically provided or offered by modeling firms. Scientific studies of hurricane formation in the Atlantic suggest a negative correlation with El Niño years, and potentially a positive correlation with La Niña years. Shear winds are one possible effect of El Niño that is thought to influence hurricane formation and development. The El Niño warmwater effect in the Pacific produces strong westerly winds at 40,000 feet that can shear the tops off Atlantic hurricanes, disrupting their formation or strengthening. In contrast, La Niña is thought to produce weak high-altitude winds that do little to prevent hurricane formation. The impact on potential insurance losses is far from clear.

Empirical studies have found no significant correlation between insurance losses and the frequency of landfall in neutral, El Niño, or La Niña years. If we assume independence of hurricane frequency and severity, the historical data indicate that the expected losses from hurricanes are lower in El Niño and higher in La Niña years. However, the available data sample is very small, and these historical observations may not be representative of the long-term trend. Researchers have also cited other variables — including African rainfall, stratospheric wind oscillation, and pressure and temperature gradients — as affecting year-to-year hurricane probabilities.[4]

- The estimation of damages is more complex for *commercial* structures than for *residential* ones. For residential buildings, there are only two major types: brick masonry and wood (apart from mobile homes, which personal line insurers typically consider a separate line of business). In the case of commercial structures, however, there is a substantial variety of construction types. They can range from small, single-proprietor stores in shopping malls to huge factories. Therefore, modeling of commercial structures is significantly more involved.
- For earthquakes, one important issue relates to the use of *time-dependent versus time-independent probabilities*. Traditional risk modeling assumes time-independent probabilities. However, as noted above, certain faults are believed to have occurrence probabilities that are time-dependent. The time-dependent model tends to decrease the current estimate of hazard in some faults while increasing it in others, based on the time since the last major event and its repeat time. For some parts of California, such as the San Francisco Bay Area, the time-dependent model would produce higher current probability estimates because the last large events happened a relatively long time ago (e.g., 1906 for the San Andreas fault, and earlier for the Hayward fault). The scientific evidence and consensus on the validity and application of time-dependent modeling is evolving and has not been constant over the last few years.[5]
- We argue later that the expected loss is considerably more important for the evaluation of a CAT bond than the entire distribution of losses. It is important to realize, however, that this does not imply that a model that correctly describes only averages of hurricane parameters or average damages to structures would be sufficient. *A model that describes these averages correctly but underestimates their variances can vastly underestimate the expected losses for CAT bonds*. This is due to the *non-linear character* of the physical processes determining the hazard intensities of the damage functions and of the financial contracts.

  For example, a hurricane of twice the average strength might do much more than twice the damage of an average strength hurricane. Similarly, both deductibles in insurance policies and the loss structure of the reinsurance contracts create non-linear relationships. This is analogous to the well-known importance of *volatility in option pricing*. The value of a call (or a put) on a stock increases if the volatility of the stock increases, even if the expected stock value at expiration of the call remains the same. Note that to be short an XOL reinsurance contract is a bear spread on the losses of the cedant — i.e., a short position in a call with one strike and a long position in a call with a higher strike.

Model Validation

There are several ways to validate and test the risk models:

- Reproduction of actual historical storms: defining the parameters of the actual event and running it against historical industry or company exposure at the time of the storm (*ex post simulations*). A similar approach is *real-time damage predictions*, wherein the modeling firm can estimate industry losses shortly after an event occurs, based on the (only approximately known) parameters of the event. This estimate can be revised as event parameter information is refined and then compared with the "actual" outcome. Unless results are analyzed across a large set of events, they are in themselves not complete because they provide a comparison only on an aggregate basis. It is important to validate the model against observed losses from various insurance portfolios that provide detailed losses by geography, by coverage, and by line of business. Furthermore, in the models, losses are stochastic, even conditional on fixed hazard parameters. Therefore, for each actual storm, we have to compare a probability distribution of losses as predicted by the model against a single number for the actual loss experience.
- Comparing *historic versus simulated loss distributions*. In this case, the frequency of different loss levels from historical events is compared with the estimated stochastic distribution.
- Comparing *measured versus predicted intensities* (wind speeds for hurricanes and ground motions for earthquakes). Using the parameters of an historic event as input to the model, have it calculate the intensities at various locations, and compare the results against measurements.
- Various *public authorities* in the states of Washington, California, Hawaii, and Florida have developed their own standards for model validations. The Florida

Commission on Hurricane Loss Projection Methodology, for example, performs an extensive validation process before it approves various models for use by insurers in Florida for rate-making purposes. This process includes various hearings, validation tests, submission of documents, and on-site visits by independent experts.
- In the process of rating CAT bond issues, the *major credit rating agencies* perform their own evaluation of the modeling. This process includes asking the modeling firms to stress-test various assumptions and mechanics in their models, hiring outside meteorological and/or geophysical experts to examine the models, and comparing relative rigor of methodology and error estimation with models used in similar structured markets, such as interest rate and prepayment modeling for mortgage securitizations.

ASYMMETRIC INFORMATION

An indemnity-based structure is attractive to issuers because it avoids basis risk, to which they would be exposed in a parametric or industry index-based deal. However, this raises various issues related to asymmetric information, such as adverse selection and moral hazard.

Adverse selection reflects the problem that the insurer or reinsurer might try to securitize the most unattractive parts of its portfolio, and keep the more profitable ones. *Moral hazard* relates to the possibility that the cedant might no longer try to limit its losses, since it has transferred the risk to investors. Risk arising from any event affecting a contract may increase from the initially estimated level if the cedant relaxes its underwriting policies (ex ante) or its claims settlement in case of a catastrophe (ex post).

Even in the absence of a relaxed underwriting policy, the risk in the covered portfolio could grow and change in composition because of the underwriting of new policies. The insurer is compensated for this *portfolio growth* by collecting more premiums. However, if the loss on the security is defined in terms of fixed attachment and exhaustion points, this could increase the risk to the securityholders without compensation.

Another issue related to asymmetric information arises out of the *proprietary nature of the exposure data*, which means the investor typically does not have full access to this information. This differs from the case of collateralized loan obligations (CLOs), where the investor might dig down and analyze each of the loans in the portfolio on his own.

There are several methods to control these issues:

- The potential for adverse selection can be addressed by including the cedant's full portfolio or complete lines of business, or by establishing in advance clear and non-negotiable rules for the selection of qualifying subject business.
- The goals of the cedant and the securityholder can be aligned by making sure that the cedant keeps a portion of all losses. First, in an XOL contract, the cedant could be required to retain a certain percentage of losses in the reinsured layer. Second, regarding losses below or above the layer, there can be limits as to how much additional reinsurance the cedant may obtain.
- Changes and growth of a portfolio can be controlled by defining a specific portfolio consisting only of pre-agreed contracts and exposures, or through periodic trigger resets or limits on the inclusion or impact of new policies in the agreement. This can also be mitigated by short risk periods, or by indexing the contract to the combined ratio (paid losses divided by premiums received) for the relevant lines of business.
- Finally, the CAT modeling firms, and to a lesser extent the bond rating agencies, play an important role in controlling the problem of asymmetric information. They have access to details of the exposure and policy data, and can provide the investor with relevant information through the probability distribution of losses. Clearly, the modeling firms then also play an important role in evaluating the quality of the data provided (completeness, consistency, etc.).

We now analyze some selected securitization transactions and discuss how these issues were addressed.

Residential Re

Exhibit 6 displays the structure of the Residential Re 1997 and 1998 deals. The term of Res Re '97 and '99 was one year, and for Res Re '98 it was fifty weeks. Losses to the bondholders are triggered if a *single hurricane* causes losses to USAA in excess of $1 billion. Participation of USAA in the reinsured layer was 20% in Res Re '97 and 10% in Res Re '98, which was set to align USAA's interests with those of the investors and thus address the issue of moral hazard. Adverse selection is prevented by including only complete lines of business of

EXHIBIT 6
Residential Re Structure

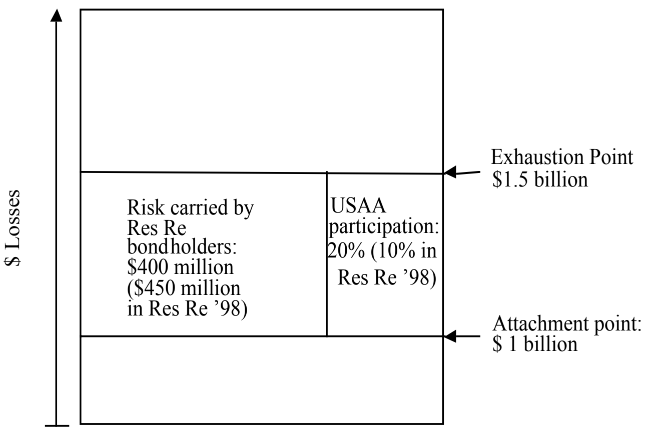

USAA (homeowners, condominium/renters, and dwelling insurance).

The issue of potential portfolio growth during the risk period is handled by limiting losses arising from new policies that are covered by the reinsurance agreement to a maximum of 9%, and by the short lengths of the risk periods.

In 1999, USAA chose to minimize their costs by optimizing across markets. USAA bought traditional reinsurance for $250 million of the $1 billion-$1.5 billion layer (i.e., 50%), issued $200 million in Res Re ’99 bonds (i.e., 40%), and, as in 1998, had to retain 10% of the risk.

Parametric Re

The Parametric Re issuance avoided many of the asymmetric information issues by use of a “parametric” structure. Two rectangular geographical grids around Tokyo were defined (see Exhibit 7), along with two schedules that relate earthquake magnitudes in these grids to losses (see Exhibit 8). If, during the risk period of the bond, an earthquake occurs within the inner grid of a magnitude 7.4, for example, 70% of the outstanding principal of the bonds is lost. Similarly, if such an earthquake were to occur within the outer grid but outside the inner grid, bondholders would lose 44% of their principal.

The parametric structure allowed this bond to have a ten-year risk period, and was particularly useful because Tokio Marine’s business covers many commercial properties, which are particularly difficult to model.

EXHIBIT 7
Inner and Outer Grids for Parametric Re

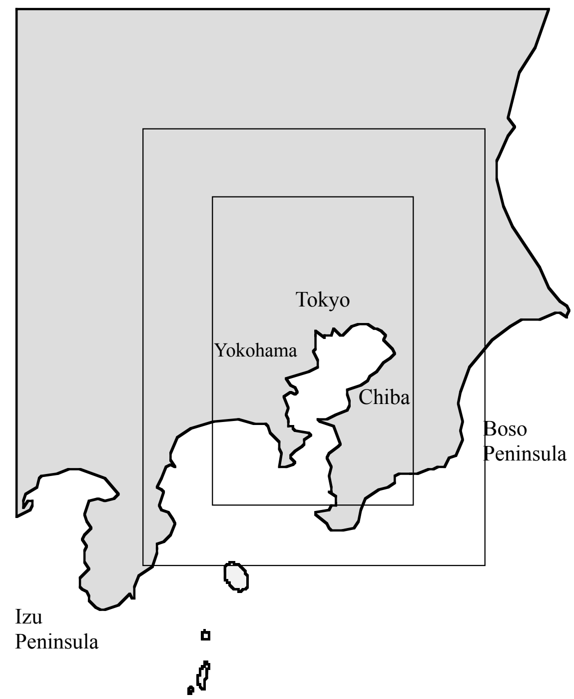

A =similar parametric structure was later used by Concentric Ltd.

Trinity Re

Exhibit 9 displays the Trinity Re ’98 structure. Portfolio growth is addressed by an innovative approach that involves the risk modeling firm RMS. When a potential loss event has occurred, the exhaustion point and the

EXHIBIT 8
Trigger Event Amounts for Parametric Re

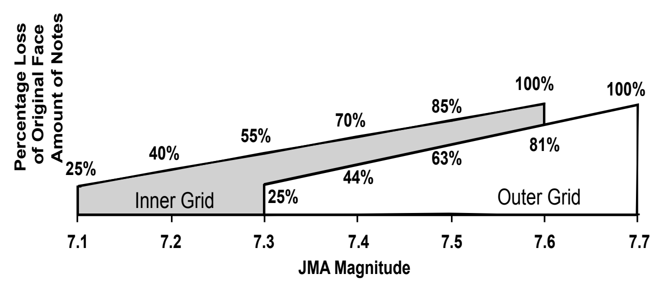

EXHIBIT 9
Trinity Re Structure

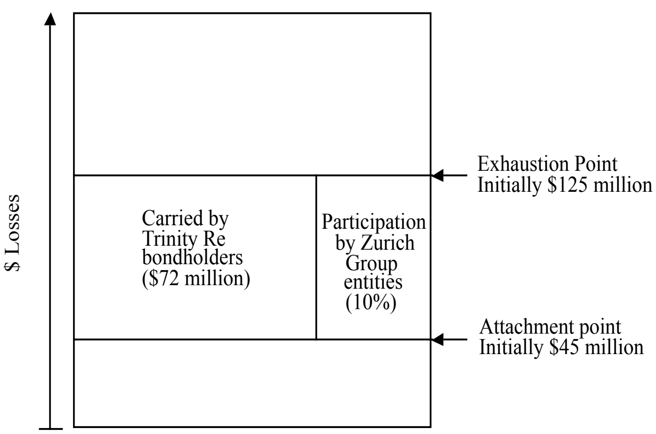

EXHIBIT 10
Mosaic Re Structure

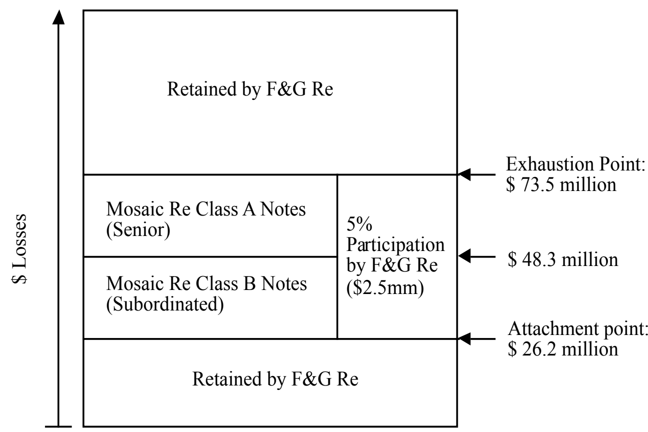

coinsurance percentage are reset using the original version of the risk model with current exposure data such that the expected annual loss approximately equals its initial value. Actual payments are then based on these reset exhaustion point and coinsurance percentages. In this way, any growth in the portfolio is accounted for, and the level of risk the bondholders are ultimately exposed to remains in a narrow range.

Similar model-based resets of trigger values are used in Trinity Re '99 and in Juno Re.

Mosaic Re

Mosaic Re I was the first "pure securitization" transaction, in that it defined a specific set of forty-two individual reinsurance contracts selected by F&G Re and put into a "defined portfolio." To the extent that there are losses with respect to this defined portfolio, F&G Re retains the first ("equity") layer of $26.2 million. Mosaic Re I security holders provide risk capital to support $45 million in excess of that retention, alongside a minimal coparticipation by F&G Re, and F&G Re again retains any losses in excess of this layer (see Exhibit 10).

An important innovation in the Mosaic Re I issuance was its division of the reinsurance contract into two risk profiles reflected by a two-class security structure. Losses between the first attachment point of $26.2 million and $48.3 million are carried by the Class B subordinated noteholders. If losses exceed $48.3 million, Class B noteholders will have lost their entire principal and risk will then be borne by Class A noteholders until the exhaustion of coverage provided by Mosaic Re at $73.5 million.

The issue of adverse selection is particularly important for Mosaic Re I, because F&G Re selected individual contracts for this securitization and other contracts to retain for its own portfolio. To reflect this, the defined portfolio was modeled on a stand-alone basis and clearly described as a specific set of risks rather than a general participation in the broader business of F&G Re. In addition, the contract provided for a high degree of risk to be retained by F&G Re. First, F&G Re kept 5% of the layers covered by Mosaic Re I Notes. Second, 100% of the first $26.2 million of losses, and 100% of any excess above $73.5 million, must also be carried by F&G Re. The retention of the first and the final layers by F&G Re, plus the coinsurance, helped to further align the interests of securityholders and the cedant.

Six months later, Mosaic Re II was issued, which had a similar structure.

SECURITIZATION AND PRICING IN THE REINSURANCE MARKETS

The transactions discussed in this article relate to a specific segment of the reinsurance market, namely *catastrophe excess of loss* reinsurance. Other segments include proportional reinsurance of property catastrophe risks and both proportional and excess of loss coverage of non-catastrophe risks, such as weather, mortgage default, casualty, life insurance, and other non-financial risks. Future securitizations may broaden to include these other classes of business. Excess of loss coverage is most frequently used in the U.S. reinsurance market, although non-U.S. markets are increasingly moving to this form of coverage because the risk assumed in XOL coverage is easier to analyze and because of the protection this coverage provides against moral hazard, adverse selection, and other risks.

As with other lines of reinsurance, the market for catastrophe excess of loss coverage is global. The U.S. is the largest market, followed by Europe and then Asia. According to a 1997 report by *sigma* (a research publication of Swiss Re), the total limits purchased and total premium paid for catastrophe excess of loss coverage in 1996 were U.S. $52.9 billion and U.S. $2.8 billion, respectively. Of the total limits purchased, the U.S. accounted for 34.5%, Europe 26.7%, and Japan 8.9%, with the remaining 29.9% was spread around other parts of the world. Of the $2.8 billion in premiums, the U.S. accounted for 46.8%, Europe for 27.4%, and Japan for 10.1%, with the remaining 15.7% from other parts of the world.

Pricing in the reinsurance markets is much less transparent than typical pricing in the financial markets, largely because there is no public and liquid market in these risks, and so only a limited amount of information is available. Many considerations influence individual reinsurance contract prices, particularly with respect to long-term relationships: The concept of reinsurance premiums paid over time as accumulating in a "bank" account implies a similar obligation on the part of the reinsured to "pay back" losses incurred. Pricing will be affected by the traditional expectation that a reinsurer will reduce prices in the absence of recent losses, by the nature of the cedant — whether it is a primary insurer (reinsurance) or another reinsurer (retrocessional coverage) — by capacity levels, and by portfolio risk concentrations.

Furthermore, comparisons across individual contracts are difficult because most contracts are indemnity-based and may have other customized features. The most transparent pricing is for "industry loss warranties," which provide for coverage based on industry loss experience and are predominantly used in the retrocessional markets. Nonetheless, it is still difficult to obtain consistent pricing information. The situation is comparable to illiquid over-the-counter options, such as out-of-the-money swaptions, where prices quoted by brokers might not be meaningful. Accordingly, our discussion regarding market demand and comparative pricing is necessarily somewhat general and qualitative.

Major U.S. catastrophes have had the biggest relative impact on the global reinsurance market, which is not surprising in view of the large percentage of worldwide property values and coverage accounted for by the United States. Hurricane Andrew in 1992 caused $18.2 billion in insured losses and led to a major dislocation of prices in the reinsurance markets.[6] This was further aggravated by the Northridge earthquake in 1994, which caused $13.5 billion in insured losses.

However, during 1993-1994, many new reinsurance companies entered the marketplace (e.g., seven new Bermuda property catastrophe reinsurers supported by $3 billion in risk capital). Many of these companies, which started with relatively small capital bases, have seen returns on equity of 15% to over 40%, so their common equity has increased substantially. Prices for catastrophe reinsurance, which increased by 50% to 200% after Andrew, have gradually come back to almost pre-Andrew levels. This cycle of falling prices and increasing equity has made it difficult to achieve revenue growth. Consequently, merger and acquisition activity has risen as the need for growth has spurred increasing consolidation in the reinsurance market.

Exceptions to this rapid price softening can be found in Florida, which is still relatively capacity-constrained. In Florida, there is a disproportionately large amount of property in high-risk locations. Consequently, primary insurers have taken on a very large exposure, while capacity constraints keep reinsurance prices relatively high. Reinsurers offering capacity in this market closely manage their risk by putting limits on their exposures both for each zone (county) and for each company.

In California, the creation of the California Earthquake Authority (CEA) in 1996 relieved severe capacity pressure and has thus led to a considerable decrease in reinsurance prices for California earthquake risk outside of the CEA in the last year. The state-sponsored CEA has been capitalized by various sources, including capital contributions from primary insurers and reinsurance. It has reduced the risk it assumes by offering homeowners policies that severely limit the amount of coverage available. For example, the policies provide for high deductibles and do not fully cover contents. The CEA offered attractive rates to carriers on its reinsurance program, considering its unprecedented size and attractive risk profile, but has recently renegotiated its reinsurance arrangements to reduce its cost of coverage. Reinsurance coverage for California earthquakes is still relatively constrained, given the existing large exposure many reinsurers have to the CEA as well as continued demand from insurers writing risks outside the CEA, such as commercial risks.

The European market was transformed by a series of major catastrophes similar to Andrew and Northridge in the U.S. in the late 1980s and early 1990s. The major events included winter storm 87J ($4.2 billion in industry losses), the series of 1990 windstorms — Daria, Vivian, Herta, and Wiebke ($11.5 billion) — and the 1988 explosion of the Piper Alpha oil platform ($2.7 billion).

Furthermore, the reconstruction and renewal process for the Lloyds market led to a contraction in catastrophe coverage from the London market, which has now reversed itself. Prices have generally remained lower in Europe relative to the peak areas of Florida and California, because more time has elapsed since the major losses, and European risk is attractive as a diversifying risk for reinsurers' portfolios. However, Europe is now moving toward the U.S. model, with a strong shift from the traditional proportional catastrophe coverage to XOL contracts.

Japan's recent major insured loss events have been Typhoon Mireille in 1991 ($6.5 billion in 1997 dollars) and the Kobe earthquake in 1995 ($2.6 billion in 1997 dollars). For earthquakes, the insurance market has traditionally been highly regulated, with very limited protection provided to homeowners and significant reinsurance by the government. Reinsurance coverage has often been purchased on a proportional basis, with an increasing shift toward XOL. With the current deregulation of the financial and insurance markets and the introduction of competition into the domestic market, demand for catastrophe reinsurance protection is expected to grow significantly.

Development of an active insurance securitization market has been driven by the capacity constraints in the traditional reinsurance markets for peak risks. In addition, the consolidation in the insurance industry has caused companies to retain more lower-level risk and seek more coverage on an XOL basis for large, severe risks. Consequently, the primary insurers that have accessed the capital markets, such as USAA and Tokio Marine, have chosen to purchase coverage at a very high "super-CAT" level.

By contrast, reinsurers have more varied reasons for accessing the capital markets. First, they tend to charge considerably more when providing reinsurance for another reinsurer (retrocessional coverage). This is because retrocessional providers are supporting the business of competitors and because high quality, detailed data are harder to acquire when analyzing a reinsurer's portfolio compared with the portfolio of a primary insurer. Furthermore, trading of risk among reinsurers only shifts risk between parties sharing similar capital bases, whose absolute number is decreasing because of mergers and acquisitions. Thus, it is increasingly attractive for reinsurers to turn to the capital markets for additional capacity.

Another important reason why insurers and reinsurers turn to the capital markets, despite the falling cost of reinsurance in the traditional market, is the issue of credit risk. It is very hard for ceding companies to tell from current reporting how much risk reinsurers have taken on, and it is an open question whether they will actually be paid by all of their reinsurers. This is a particularly serious issue following a major catastrophe, when a large number of insurers will be turning to reinsurers for payment all at the same time. After Andrew, a significant number of insurance and reinsurance companies were simply declared insolvent. An insurance-linked security presents no credit risk, as the limit is typically fully funded by the proceeds of the issuance held in trust.

Other forces supporting the development of the insurance securitization market include diversification of risk capital sources as a hedge against the historical volatility in reinsurance pricing and capacity after major events, attractive product features such as fixed-price multiyear coverage, and — most important — the competitive advantage that less-expensive fixed-income risk capital can provide to support individual insurers and reinsurers that tap capital markets to expand their business more quickly.

RELATIVE VALUE

In this section, we analyze the pricing and relative value of CAT bonds. First, we compare a single CAT bond with a single high-yield bond. We use a binomial one-period model to analyze their mean-variance profiles, and then look at the whole probability distribution. Second, we analyze the impact of inter- and intra-asset class diversification. Finally, we briefly note other valuation considerations.

Single CAT Bond versus Single High-Yield Bond in a One-Period Setting

Mean-Variance Analysis. We start the relative value analysis in the framework of a simple binomial model. This model allows us to compare high yield bonds and CAT bonds within the same framework. We assume that an investor buys the bond at the beginning of the risk period at par ($100). At the end of the risk period, the investor receives a stochastic dollar amount $\tilde{V}$. With probability p, a default event or catastrophe has occurred, in which case the investor receives a recovery value $\tilde{R}$, which itself is a stochastic variable. With probability (1 – p), no event has occurred, in which case the investor receives 100 + r + s, i.e., the principal, plus the risk-free interest over the period, plus a promised spread (see Exhibit 11).

We assume that r and s are non-stochastic. Furthermore, we ignore issues of multiple cash flows and the applicable reinvestment rate. Therefore, this model is appropriate for securities with a short risk period, such as Residential Re, Trinity Re, or Mosaic. We also apply it to Parametric Re, which is a ten-year security. In this case, we assume Parametric Re consists of ten independent one-year deals, and apply the model to the first year. We return to the issue of multiyear deals in a discussion of the limitations of the one-period analysis.

For Trinity Re, the risk period is ten months, and for Residential Re '98 it is fifty weeks. In these cases, r and s have to be taken as the risk-free return and promised spread *over the risk period* (rather than annualized).

If we denote the expected recovery rate by $E[\tilde{R}]$ and its variance by $var(\tilde{R})$, we have:

$$E[\tilde{V}] = (1 - p)(100 + r + s) + pE[\tilde{R}] \qquad (1)$$

$$var[\tilde{V}] = (1 - p)\{(100 + r + s) - E[\tilde{V}]\}^2 + p\{var[\tilde{R}] + (E[\tilde{R}] - E[\tilde{V}])^2\} \qquad (2)$$

These formulas follow from the law of iterated expectations (see the appendix).

We assume that the risk-free rate for one year (U.S. Treasury rate) is r = 5.50%, and that the swap spread between Treasuries and LIBOR is 0.40%. Thus, the spread is s = 0.40% + (spread over LIBOR). Results can be found in Exhibit 12.[7]

Various comments are in order:

- The *default probabilities for the various high-yield grades* are taken from a 1998 Moody's report (see "Historical Default Rates of Corporate Bond Issuers" [1998]). The numbers we show are from Exhibit 15 of the Moody's report; they refer to the observed one-year default rates by rating in the period 1983-1997. Keep in mind that these observed default probabilities relate to a relatively small sample size, and thus statistical fluctuations are possible. In particular, for the Ba2 category, it can be argued that the default probability of 0.60% is somewhat out of line when compared with other default probabilities, and that the long-term average should be higher. Furthermore, the economic environment is not static. At any given time, market participants may believe that default rates over the next year will be higher or lower than their long-term averages. In particular, we find negative expected excess returns and Sharpe ratios for the B2 category. This might indicate that market participants believe that B2 default rates will be lower than 6.70%.
- We are considering bonds that have maturities of one year. *Spreads for one-year high-yield bonds* are difficult to estimate, because these bonds are not very actively traded. The spreads given in Exhibit 12 indicate our estimates at the time the table was constructed. They were based on Bloomberg's fair market curves for industrial high-yield bonds, which we have modified (increased) to account for estimates by Goldman Sachs traders. We emphasize again that we are considering bonds with one year until maturity. High-yield bonds with longer maturities have higher spreads and therefore higher expected returns. However, they also have higher risks (standard deviations), because of the additional spread risk. In the case of no default, the dollar return at the end of the period becomes a random variable, which can be higher or lower than 100 + r + s.
- For corporate bonds, we use recovery rates and standard deviations thereof as given for *senior unsecured debt* in Exhibit 25 of the Moody's report.

EXHIBIT 11
A Binomial Model for Relative Value Analysis

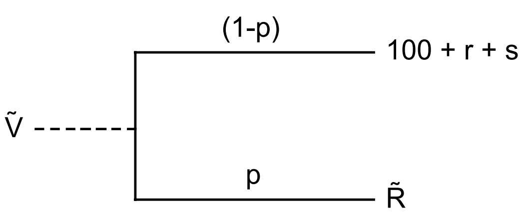

EXHIBIT 12

Relative Value Analysis*

Speculative-Grade	Historical (1983–1997) Default Probabilities	Spread Over LIBOR	Recovery Rate (%)	Std Dev of Recovery	Std Dev of Return	Expected Loss	Sharpe Ratio
	p		$E(\tilde{R})$	$SD(\tilde{R})$	$SD(\tilde{V})$		
Ba2	0.60%	1.10%	51.26	25.81	4.75	0.33%	0.25
Ba3	2.70%	1.36%	51.26	25.81	10.02	1.51%	0.02
B1	3.80%	1.84%	51.26	25.81	11.91	2.15%	0.01
B2	6.70%	2.00%	51.26	25.81	15.66	3.79%	-0.09
B3	13.20%	2.49%	51.26	25.81	21.49	7.54%	-0.22
Principal at Risk Tranche	**Attachment Probs**						
Res Re '97	1.00%	5.82%	48.30	30.60	7.01	0.63%	0.80
Parametric	1.02%	4.36%	41.23	30.04	7.57	0.70%	0.54
Trinity '98	1.53%	3.67%	54.61	38.27	8.14	0.83%	0.39
Res Re '98	0.87%	4.04%	42.67	35.72	7.06	0.58%	0.54
Mosaic I Class A	1.13%	4.40%	61.40	30.05	6.06	0.55%	0.70
Mosaic I Class B	4.29%	8.20%	52.98	32.71	14.09	2.62%	0.42
Principal-Protected Tranche	**Attachment Probs**						
Res Re '97	1.00%	2.76%	75.05	16.22	3.72	0.34%	0.76
Parametric	1.02%	2.09%	73.47	15.02	3.78	0.35%	0.56
Trinity '98	1.53%	1.57%	80.91	18.14	3.86	0.39%	0.39
Mosaic I	1.13%	2.15%	83.53	15.03	3.03	0.28%	0.75

*For CAT bonds, we have multiplied the quoted spreads by #d/360, where #d is the total number of days over which interest is paid (see the appendix).

- We define the Sharpe ratio as the ratio of the expected excess return (over the risk-free rate) to the standard deviation of the dollar return. Both numbers relate to the risk period: one year for Res Re '97, Mosaic, and the high-yield bonds; ten months for Trinity; and fifty weeks for Mosaic. For Parametric, we also assume a one-year risk period. We have *not annualized* these numbers (except for Parametric).
- We use the *Treasury rate*, not LIBOR, as the *risk-free rate*. This increases our expected excess returns and thus our Sharpe ratios.
- For the CAT bonds, the attachment probabilities, expected recovery rates, standard deviations of recoveries, and expected losses are either taken directly or calculated from *information contained in the offering circulars*. The *recovery values include the interest payments*, which will be paid even in the case of a total loss of principal.
- The spreads to LIBOR for the CAT bonds denote the *issuance spreads*. For deals with a risk period of less than one year (Trinity Re, Res Re '98, and Mosaic), these spreads, as well as the attachment probabilities and the expected losses, *relate to the actual risk periods; they are not annualized*.

We find that in terms of the Sharpe ratios, the CAT bonds look much more attractive than the high-yield bonds. This large difference does not disappear if we vary the high-yield default probabilities or their spreads within ranges that look reasonable to us. For example, future default rates for the Ba3, B1, and B2 categories would need to be about 10% of their historical averages to give them Sharpe ratios around 0.5.

Of course, there are other important issues in the valuation of securities, beyond their individual means and variances over a single period. First, these distributions are very far from normal distributions, and thus higher moments (for example skewness) might be important to investors. Second, the price process for CAT bonds is very different from that of high-yield bonds, in that the CAT bond has a larger Poisson component (jump process). We analyze these and other issues in the following subsections.

Stochastic Dominance. In Exhibits 13 and 14, we compare the probability distributions of returns for Res Re '98 and Mosaic with those of high-yield bonds. For the recovery value distributions of high-yield bonds, we have assumed a beta distribution with mean ($51.26) and standard deviation ($25.81) as found by Moody's. For the distributions of the CAT bonds, we use the distributions calculated by the modeling firms.

Let us briefly define the concept of (first-degree) *stochastic dominance*. If the probability of asset A's rate of return exceeding *any* given level is larger than or equal to that of asset B's rate of return exceeding the same level, then asset A stochastically dominates asset B. Everything else being equal (in particular correlations of asset A and B to other assets), asset A is preferred by any rational investor.

Under the assumptions we made, we find that Residential Re '98 stochastically dominates B3- and maybe B2-rated bonds, and that Mosaic dominates B2 and maybe B1. Suppose we are looking for a one-period buy-and-hold investment and can choose only between Res Re '98 and a B3-rated bond. Then the Res Re '98 bond looks more attractive than the B3 bond, independent of our risk aversion.

If one believes that the catastrophe bonds stochastically dominate the high-yield bonds, this is a much stronger statement than to say they have a larger Sharpe ratio. However, the statement of stochastic dominance is more sensitive to assumptions, such as default probabilities and in particular conditional distributions of recovery values for high-yield bonds.

Limitations of the One-Period Analysis. For multiyear bonds, such as Parametric Re, several issues that we have ignored in this analysis could become more relevant. Spreads in the CAT bond market can widen or tighten leading to price movements in outstanding securities. In addition, we consider both the risk-free rate and the swap spread between Treasury and LIBOR as constants. Over longer periods, the stochastic nature of these rates becomes more important, as does the timing of the various cash flows. For high-yield bonds, features such as calls at a premium can also affect the total return. In a complete analysis of multiyear CAT bonds, these issues also need to be addressed.

Another difference between high-yield bonds (or emerging market bonds) and CAT bonds lies in the dif-

EXHIBIT 13
Stochastic Dominance: Res Re '98

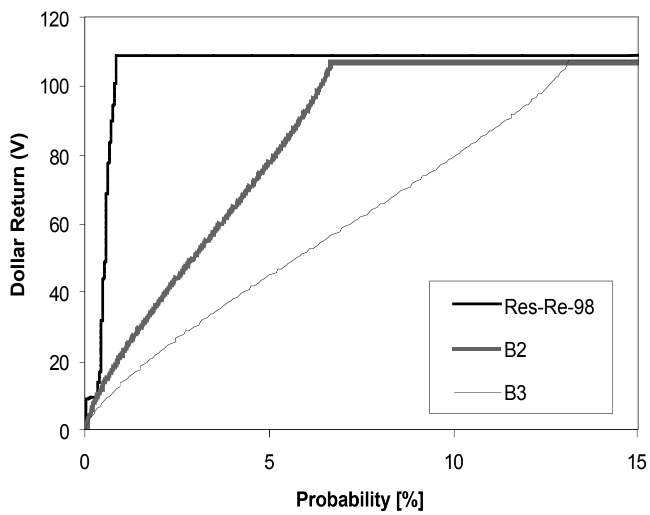

This graph shows, e.g., that at a tail probability of 5%, the dollar return of a B2 bond will exceed $78 with a confidence level of 95%.

EXHIBIT 14
Stochastic Dominance: Mosaic Class A Notes

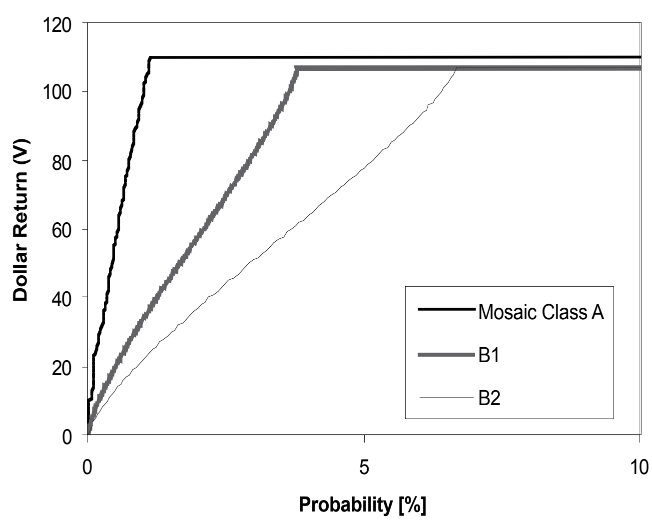

ferent information flow and therefore price movement process. In the case of high-yield bonds, new information about the financial condition of the issuer tends to arrive gradually. A severe hurricane event would strike with only a few hours notice.[8] Even more extreme, an earthquake strikes with no advance warning at all.

For the price of the security as a function of time, this implies that high-yield bond prices have a larger diffusion component (Brownian motion), while CAT bond prices are characterized by large sudden jumps (jump process). Investors may consider this a drawback of CAT bonds compared with high-yield (and emerging market) bonds, because they feel that they might be able to limit their losses by selling the bonds before an actual default.[9] Consequently, they might require a premium for CAT bonds to compensate for this difference.[10]

Intra- and Interasset Class Diversification

Intra-Asset Class Diversification. An important difference between CAT bonds and high-yield or emerging market bonds as an asset class lies in the large additional potential of CAT bonds for both intra- and interasset class diversification. Diversification potential within high-yield bonds is limited because price and spread behavior is generally correlated with conditions in the "credit markets." No matter how many high-yield bonds an investor holds, if macroeconomic credit concerns cause spreads to widen, then the portfolio will perform poorly.

The new market for insurance-linked securities offers a broad range of independent and uncorrelated non-financial risks. If an investor holds a single CAT bond, there is a large probability of a high return, but a small probability of losing the entire principal. However, if the investor can diversify across several CAT bonds with independent risks, the risk of losing the entire principal of that aggregated holding becomes negligible.

We illustrate this in Exhibits 15 and 16 for hypothetical CAT bonds that pay $110 at maturity with a 99% probability, but only the coupon of $10 with 1% probability. It can be seen that with ten independent risks, a dramatic improvement can be achieved through diversification. Observe that the probability of getting a return of less than –10% is virtually nil in the diversified portfolio.

We have simplified this analysis by considering ten binary CAT bonds with identical independent return distributions, but the main result regarding the enormous benefit of diversification within CAT bonds applies equally to actual CAT bonds. In this context, it might be interesting to mention what different natural catastrophe risks exist that are immediate targets for securitization:

1. Southeast U.S. (mainly Florida) hurricane.
2. Northeast U.S. (Long Island, New York) hurricane (very rare, but potentially causing very large insured damages).
3. Northern California earthquake.
4. Southern California earthquake.
5. New Madrid earthquake.
6. Japan earthquake.
7. Japan typhoon.
8. Canada earthquake.
9. Europe windstorm, flood, hail.
10. Europe earthquake.
11. Israeli earthquake.
12. New Zealand earthquake.
13. Australia earthquake.
14. Australia typhoon.

EXHIBIT 15

Return Distribution for a Hypothetical (Binary) CAT Bond

Probability	Return
99%	10%
1%	–90%

EXHIBIT 16

Return Distribution for Ten Independent CAT Bonds of the Same Type as in Exhibit 15

Probability	Return
90.44%	10%
9.14%	0%
0.42%	–10%
1.1×10^{-4}	–20%
2.0×10^{-6}	–30%
2.4×10^{-8}	–40%
2.0×10^{-10}	–50%
1.2×10^{-12}	–60%
4.4×10^{-15}	–70%
1.0×10^{-18}	–80%
1.0×10^{-20}	–90%

Beyond natural catastrophe risks, certain corporate risks share the same characteristic of low frequency, high severity. Examples are catastrophic events at commercial airlines, oil platforms, or nuclear reactors.

Another potentially huge area for diversification is weather risks. These can relate to temperature (heating or cooling degree days) or to levels of precipitation. Although they may not share the low-frequency, high-severity characteristic, they also have inherently low correlation with financial risks. Other insurance-related risks that have been packaged and sold in a similar fashion include the residual value risk in auto leases, mortality risk (life insurance), and mortgage default risk.

The Market Model. Because insurance-linked securities are uncorrelated with other financial markets, their market risk is zero. Let us consider the market model below:

$$\tilde{R} = \alpha + \beta\tilde{M} + \tilde{e} \qquad (3)$$

Here, $\tilde{R}$ denotes the stochastic return of a portfolio, $\tilde{M}$ is the stochastic return of the market, α denotes the "abnormal" excess return, and $\tilde{e}$ is the idiosyncratic risk. The market component of the risk, $\beta\tilde{M}$ cannot be diversified away, while the idiosyncratic risk can be diversified away and thus does not command excess return.

This relation is displayed graphically in Exhibit 17. cos(δ) is the correlation between the return of the portfolio and the market return. If the angle is zero, the cosine is one and the correlation is one. If the angle is 90 degrees, then the cosine — i.e., the correlation — is zero.

Now, if we can find assets that:

1. contain excess return, and
2. are orthogonal to the market, and if

EXHIBIT 17
Market versus Idiosyncratic Risk

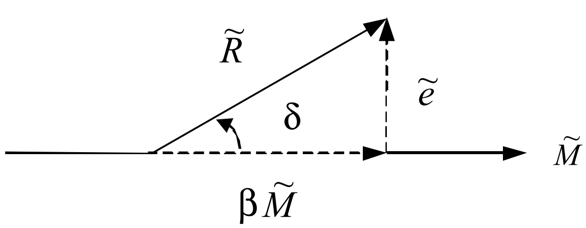

EXHIBIT 18
Enhancing Performance by Adding High-Alpha Assets

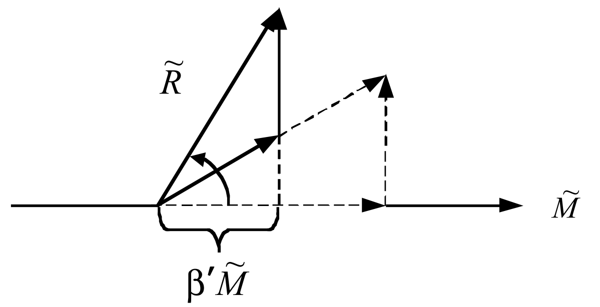

3. we add those assets to the portfolio, then
4. the market component of the total risk of the portfolio decreases, and
5. the α of the portfolio increases.

This is shown in Exhibit 18.

If the market has to be tracked, a total return swap can be used. In a total return swap, the investor pays LIBOR plus a small spread, and receives the total return on a market index. This type of instrument allows an investor to manage the risk profile by adjusting the beta of the portfolio back to the level desired, as shown in Exhibit 19.

Insurance-Linked Securities in a Diversified Portfolio. In this section, we analyze numerically the impact of moving a small part of a large portfolio into one or several CAT bonds. We look at the change in the expected excess return and in the standard deviation, and we exam-

EXHIBIT 19
Recovering Beta Using a Total Return Swap

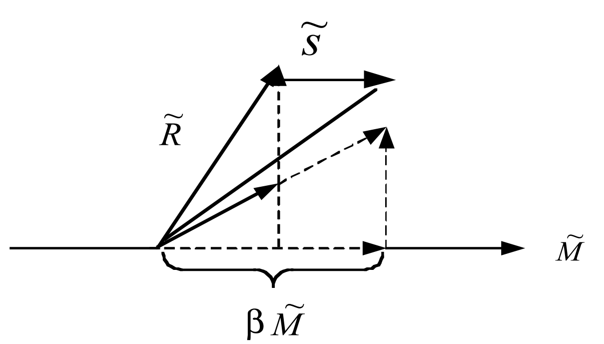

ine the impact of the investment on higher moments of the return distribution. Assume we start with a large diversified portfolio of bonds with an expected excess return of 1.7% and a standard deviation of its return of 9.2%. This might represent a bond portfolio that includes a high-yield or emerging market component. (Our analysis can easily be generalized to any other assumptions for the expected return and standard deviation.)

We now assume that we sell a percentage of the portfolio and buy a single CAT bond from the proceeds, either Mosaic Class A or Class B (at its issuance price and spread). The expected excess return increases, as Exhibit 20 shows.

We display the standard deviation of the resulting portfolio in Exhibit 21 and the Sharpe ratio in Exhibit 22. Note that the increase in expected return is always larger for Mosaic Class B than for Mosaic Class A. If only a small percentage of the portfolio is invested in the security, however, the standard deviation decreases by the same amount. The higher relative risk of the Mosaic Class B matters only if the percentage invested in the new security exceeds 5%-7%. Therefore, the Sharpe ratio of the portfolio is larger if we choose to invest in the riskier Class B rather than in Class A.

The intuition for these results lies in the fact that both securities are totally uncorrelated with the remainder of the bond portfolio. As we prove in the appendix, if the component invested in the CAT bond is small, then the standard deviation of the new portfolio is dominated by the original portfolio bonds, and the contribution of the new CAT bond to the total standard deviation is highly suppressed.

This remains true even if the correlation of the new asset to the original portfolio is small, but not exactly zero. This can be seen in Exhibit 23, where we consider moving a percentage of the diversified portfolio into a new asset that has the return and standard deviation of Mosaic Class A, but a range of correlations with the bond portfolio. For a small correlation such as 5%, and a small component in the new asset, the results are essentially identical to those with zero correlation. We prove this algebraically in the appendix.

We can decrease the risk of the portfolio further by diversifying across several non-financial risks. Consider the amount we have in each subset of insurance-linked securities with independent risks. The first subset might contain several Florida hurricane CAT bonds, the second subset one or more California earthquake bonds, and so on.

EXHIBIT 20

Expected Excess Return of the Portfolio as a Function of the Percentage Moved into Mosaic

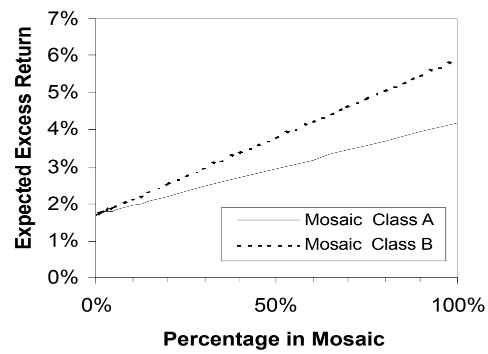

EXHIBIT 21

Standard Deviation of the Portfolio as a Function of the Percentage Moved into Mosaic

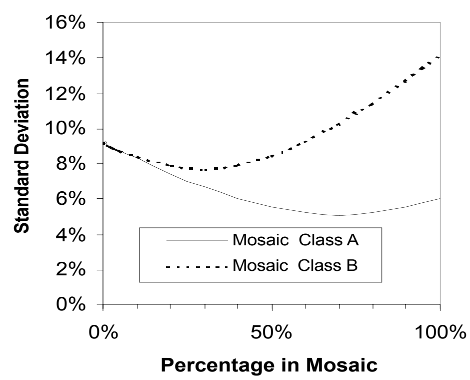

EXHIBIT 22

Sharpe Ratio of the Portfolio as a Function of the Percentage Moved into Mosaic

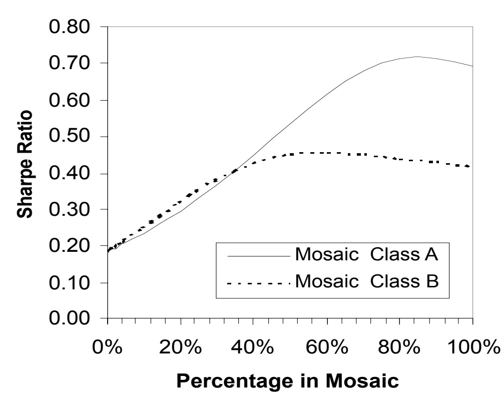

EXHIBIT 23

Impact of the Correlation Between the New Asset and the Original Portfolio on the Standard Deviation of the New Portfolio

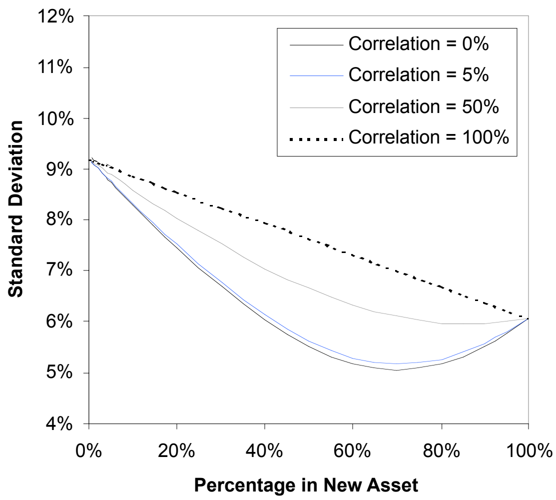

We assume that the amount we have in each independent risk category is small compared with the amount still invested in the original diversified portfolio, so that the contribution of the insurance-linked securities to the standard deviation (i.e., to the risk) of the total portfolio is much smaller than their aggregate contribution to the expected return. Furthermore, their contribution to the skewness and kurtosis of the portfolio is even smaller, as is their contribution to all higher moments.

In other words: Under the assumptions that we have made, the form of the return distribution is very close to that of the original portfolio, even if the CAT bonds have a very different (typically rather binary) return distribution.

Consider an illustrative example. Suppose we move 20% of a diversified portfolio into CAT bonds. If we put 5% into a Japan earthquake (JP EQ) bond, 5% into a Japan typhoon bond (JP TY), and 10% into several Florida hurricane (FL HR) bonds, we find:

Expected return of new portfolio =
80% × (expected return of market portfolio) +
5% × (expected return of JP EQ bond) +
5% × (expected return of JP TY bond) +
10% × (expected return of FL HR bonds)

However, for the variance of the new portfolio, we have:

Variance of new portfolio =
[0.8^2 × (variance of market portfolio)] +
[0.05^2 × (variance of JP EQ bond)] +
[0.05^2 × (variance of JP TY bond)] +
[0.10^2 × (variance of FL HR bonds)]

= 0.64 × (variance of market portfolio) +
0.0025 × (variance of JP EQ bond) +
0.0025 × (variance of JP TY bond) +
0.01 × (variance of FL HR bonds)

It can be seen how the contribution of the CAT bonds to the variance is strongly suppressed. This is even more so for the third (central) moment (skewness) and the fourth moment (kurtosis):

Third moment of new portfolio =
[0.8^3 × (third moment of market portfolio)] +
[0.05^3 × (third moment of JP EQ bond)] +
[0.05^3 × (third moment of JP TY bond)] +
[0.10^3 × (third moment of FL HR bonds)]

= 0.512 × (third moment of market portfolio) +
0.000125 × (third moment of JP EQ bond) +
0.000125 × (third moment of JP TY bond) +
0.001 × (third moment of FL HR bonds)

and so on. So the skewness, kurtosis, and higher moments of the new portfolio are essentially determined by the market portfolio. This remains true even if the CAT bonds are not exactly independent from each other and the market, as long as the relevant correlation coefficients are small. We prove the following results more rigorously in the appendix:

1. If we move only a very small amount into independent new assets, we need to consider *only* the expected return of these new assets; neither their variance nor their higher moments matter.
2. If the expected return of these new assets is greater than the risk-free rate, then the Sharpe ratio of the portfolio can be increased by moving a small amount into the new assets.
3. If the original portfolio has a normal return distribution, then so does the new portfolio, as long as the amount in the CAT bonds is small. To a very good approximation, the distribution of the CAT bonds

does not matter.

Note that this justifies the use of the mean-variance approach when adding small CAT bond components to a diversified portfolio that has an approximately normal return distribution, although the distribution of the CAT bonds is far from being normal.

Exhibit 24 shows the results of an analysis of the moments of the distribution of returns for portfolios with CAT bonds.[11] We assume that the return distribution of the original portfolio is normal. We then consider the availability of either one or five independent CAT bonds, and move increasing percentages of the portfolio into these (equal percentages into each CAT bond). We have modeled the CAT bonds after Res Re '98, with a spread over LIBOR of 400 bp, but this assumption is not essential. It is remarkable that a relatively large component of the portfolio can be moved into insurance-linked securities before the deviation from the normal distribution becomes noticeable. With a single CAT bond available, 10% to 20% can be moved into the CAT bond without distorting the distribution very much, and with five independent CAT bonds, the CAT bond component could be as high as 50%. This is shown graphically in Exhibits 25 and 26. This illustrates point 3 above, which we prove in the appendix.

Notice how much the expected return can be increased and the risk decreased by adding CAT bond exposure to the bond portfolio. Moving 5% into a single CAT bond increases the average return by 11 bp while

EXHIBIT 24

Moments Analysis of Portfolios with CAT Bonds

Number of Independent CAT Bonds	Total % in CAT Bonds	Mean Excess Return	Standard Deviation of Return	Sharpe Ratio	Skewness	Kurtosis	Normalized Fifth and Sixth Central Moments	
0	0	1.70	9.20	0.18	0.00	3.00	0.00	15.00
1	1	1.72	9.11	0.19	0.00	3.00	0.00	15.00
1	2	1.74	9.02	0.19	0.00	3.00	0.00	14.99
1	5	1.81	8.75	0.21	0.00	2.99	0.00	14.93
1	10	1.91	8.31	0.23	–0.01	2.97	–0.01	14.69
1	20	2.12	7.49	0.28	–0.09	3.01	–0.56	14.91
1	30	2.34	6.78	0.34	–0.40	4.10	–7.05	41.28
1	40	2.55	6.20	0.41	–1.23	9.38	–46.45	297.87
1	50	2.76	5.80	0.48	–2.93	25.10	–198.20	1,660.16
1	60	2.97	5.61	0.53	–5.58	57.10	–581.07	6,023.05
1	70	3.18	5.66	0.56	–8.63	101.33	–1,202.26	14,409.76
1	80	3.40	5.94	0.57	–11.15	142.32	–1,841.75	24,039.76
1	90	3.61	6.42	0.56	–12.57	167.02	–2,250.20	30,571.66
1	100	3.82	7.06	0.54	–12.97	174.10	–2,370.00	32,535.00
5	1	1.72	9.11	0.19	0.00	3.00	0.00	15.00
5	2	1.74	9.02	0.19	0.00	3.00	0.00	15.00
5	5	1.81	8.74	0.21	0.00	3.00	0.00	14.99
5	10	1.91	8.29	0.23	0.00	2.99	0.00	14.93
5	20	2.12	7.39	0.29	0.00	2.96	0.00	14.67
5	30	2.34	6.51	0.36	–0.02	2.89	–0.01	14.08
5	40	2.55	5.66	0.45	–0.06	2.80	–0.12	13.03
5	50	2.76	4.86	0.57	–0.20	2.79	–0.76	12.26
5	60	2.97	4.14	0.72	–0.56	3.40	–4.26	19.37
5	70	3.18	3.54	0.90	–1.42	6.43	–20.23	81.01
5	80	3.40	3.12	1.09	–3.06	15.22	–73.13	363.51
5	90	3.61	2.99	1.21	–4.99	28.55	–165.22	965.03
5	100	3.82	3.16	1.21	–5.80	34.82	–211.98	1,301.40

EXHIBIT 25

Return Distribution for a Portfolio with 10% in a Single CAT Bond and the Remaining 90% in a Diversified Bond Portfolio, for Which We Assume a Normal Return Distribution

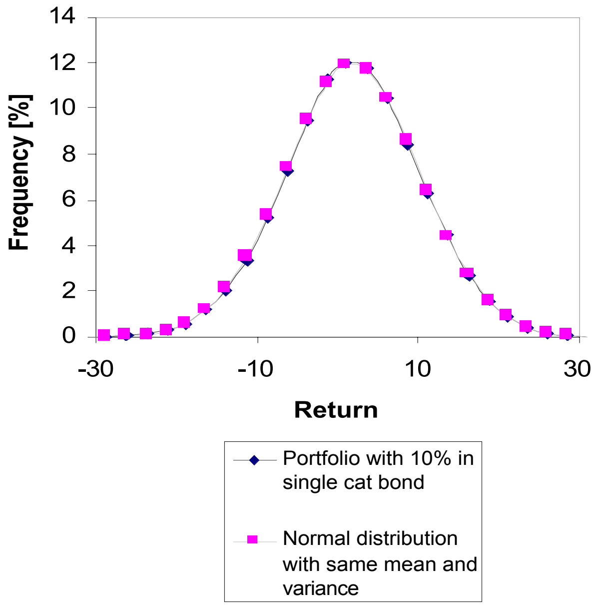

EXHIBIT 26

Return Distribution for a Portfolio with 50% in Five Independent CAT Bonds and the Remaining 50% in a Diversified Bond Portfolio, for Which We Assume a Normal Return Distribution

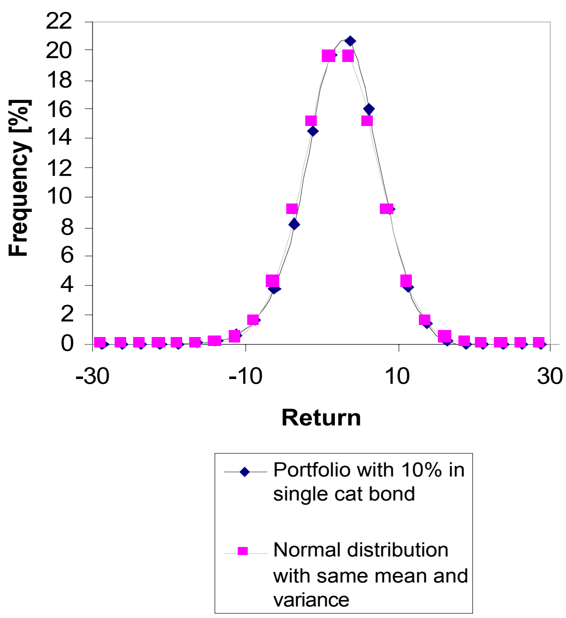

simultaneously decreasing the risk (standard deviation) from 9.20% to 8.75%. Moving 10% into five independent but similar CAT bonds increases expected return by 19 bp, but decreases the risk further to 8.29%.

Seasonality of Hurricane Bonds

While earthquakes can strike at any time during the year with equal probability, hurricanes exhibit strong seasonality. This can be seen from the historical distribution displayed in Exhibit 27. Most hurricanes make landfall during July through November, with a peak in August and September.

This implies that if no serious hurricane has happened by the end of September, the risk for a hurricane bond will be greatly reduced by then. Now, how would we expect this seasonality to affect the pricing of a hurricane bond throughout the year?

We look at a hypothetical bond with a risk period from June 15 of year 1 through June 15 of the next year. We assume it is issued at a spread of 400 bp over the one-year U.S. Treasury and has an expected loss of 60 bp:

T = maturity = 1 year
$s(t = 0) = 400$ bp
$E(L \mid t = 0) = 60$ bp

Now to find the remaining expected loss at a later time, we integrate the normalized annual distribution of

EXHIBIT 27

Seasonal Distribution (Probability Density as a Function of Time) of Landfalling Hurricanes, Smoothed Historical Data, 1949–1997

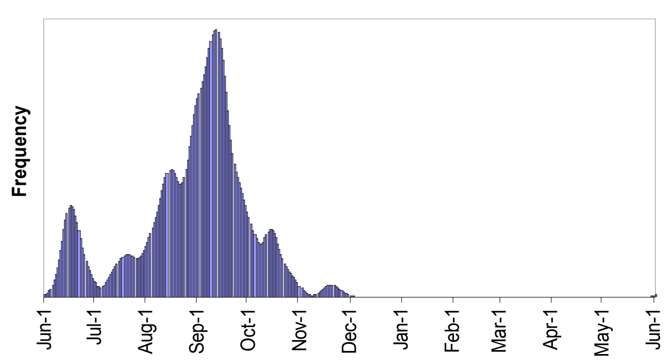

We have smoothed the data by smearing any single event using Gaussian distribution with a standard deviation of five days.

EXHIBIT 28

Cumulative Probability Distribution Function for Hurricane Dates, Derived from the Data in Exhibit 27

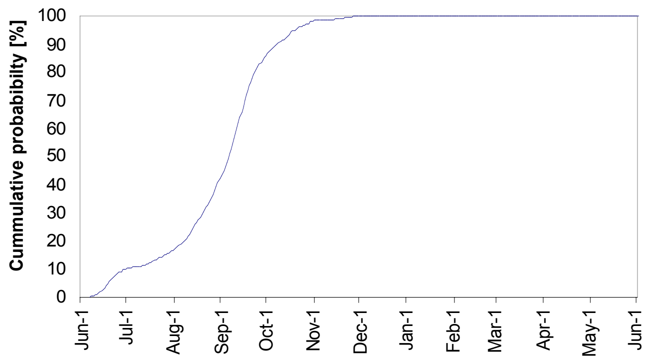

EXHIBIT 29

Spread as a Function of Time for a Hurricane Bond, Under the Assumptions Described in the Text

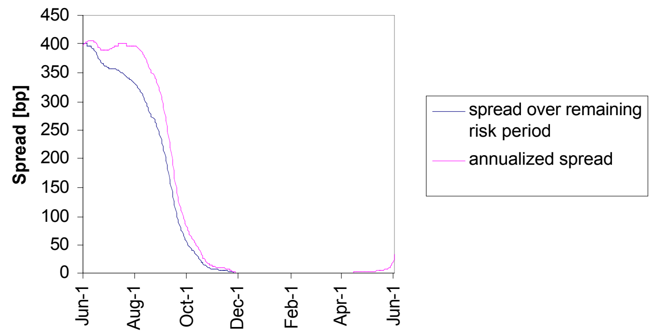

hurricanes in Exhibit 27, starting at the issue date June 15. The result R(t), which is displayed in Exhibit 28, indicates how much of the risk is gone at any later point in time.

Therefore, the expected loss at a later point in time is reduced correspondingly:

$$E(L\,|\,t) = [1 - R(t)]E(L\,|\,0)$$

Now *if we assume that the spread over the one-year Treasury is a constant multiple of the expected loss*, then the spread *over the remaining risk period* at a later time is given by:

$$s(t) = [1 - R(t)]s(0) \tag{4}$$

and the annualized spread is:

$$s_{annualized}(t) = \frac{1 - R(t)}{T - t}s(0) \tag{5}$$

We plot the result in Exhibit 29. Note, however, that we have made the following assumptions:

- We assume that no catastrophic hurricane has happened until time t, which would affect the principal of the bond.
- We assume that the risk multiple stays constant. Actually, however, risk multiples tend to increase for risks with very small probability and decrease for risks with higher frequency. (The ratio of the spread over Treasuries to the expected loss from defaults is much greater for investment-grade corporates, such as AAA or AA bonds, than it is for high-yield bonds.) Therefore, we could expect that where we indicate very small spreads in Exhibit 29, the market might actually trade at higher spreads.
- We assume that spreads do not move because of general changes in the market, such as changing perceptions of CAT bonds or movements of spreads in other markets.

In Exhibit 30, we display the price of the CAT bond under the same assumptions, based on:

$$P(t) = \frac{1 + (T - t)[r_f + s(0)]}{1 + (T - t)[r_f + s_{annualized}(t)]} \tag{6}$$

EXHIBIT 30

Price of the Hurricane Bond as a Function of Time

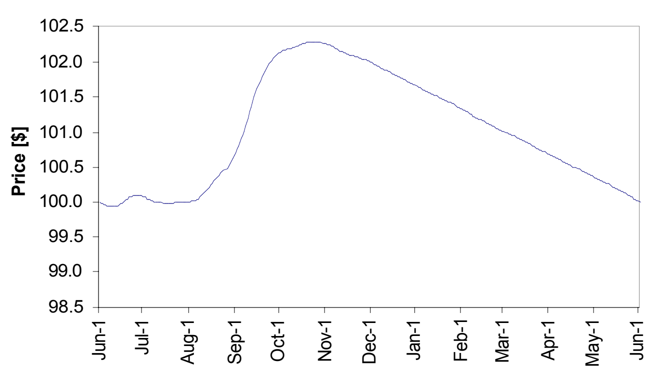

Other Valuation Considerations

One issue that is becoming increasingly important with the growth of the CAT bond market is that of correlations between different insurance-linked securities. Worldwide, we might be able to find ten to twenty completely independent catastrophe risks that are suitable for securitization. However, many of the recently issued CAT bonds have a large Florida hurricane exposure, and therefore have some degree of positive correlation.

On the other hand, they have exposure concentrations in different parts of Florida, and to different underlying insurers. With the current level of disclosure, a detailed analysis of these correlations is difficult. Greater detail of underlying exposure — in the form of aggregated policy information, modeling firm event output files, or some other comparable granular form that would allow comparison of relative exposure and potential correlation estimation — would be a very attractive feature for future transactions and aid greatly in trading and expanding further the popular acceptance of these instruments.

Another consideration is liquidity, which might be smaller for CAT bonds than for high-yield bonds. As we have mentioned before, this issue is particularly important for multiyear CAT bonds such as Parametric Re.

Finally, the risk in insurance-linked securities is somewhat different from the risk in other financial assets, whose analysis is familiar to investors. So there is an additional cost of research for investors, and they have to find their level of comfort with the risk assessment provided by the modeling agencies.

SUMMARY AND OUTLOOK

We analyze the relative value of securities that are linked to natural catastrophes. We develop a framework for their evaluation, and in particular we show how uncorrelated new assets with high excess returns can be used to enhance the return and the Sharpe ratios of diversified portfolios. Moving a portion of a diversified portfolio into a number of different independent risks offers a significant potential for yield enhancement and simultaneously for risk reduction.

We have concentrated on natural catastrophe bonds, but most of our conclusions are valid for any other asset with low correlation with existing instruments. This relates to other types of insurance-linked securities, such as weather-related instruments, automobile lease residual value securities, and securitization of mortality risk for life insurers. Indeed, the greater the number of available independent risks, the larger the potential for portfolio yield enhancement and risk reduction.

With repeat issuers, an increasing number of new CAT-related issues, and other innovative securitizations, the market for insurance-linked securities is developing and growing. We believe this will continue to provide issuers with access to additional high credit quality capital at reasonable prices, and investors with securities with attractive and unique risk/reward profiles.

APPENDIX

Insurance-Linked Securities in a Diversified Portfolio: Analytics

We start with a portfolio Π. We assume the expected return and its standard deviation over some fixed time period are:

$$E\{\Pi\} = M \qquad SD\{\Pi\} = S \tag{7}$$

and its higher central moments are:

$$M_{N\text{-th}}^{(central)} = E\{[\Pi - M]^N\} \tag{8}$$

Now assume there is another asset V, into which we move part of the portfolio. This new asset might be a CAT bond or some other insurance-linked security. Its mean and standard deviation are denoted by:

$$E\{V\} = \mu \qquad SD\{V\} = \sigma \tag{9}$$

and we are interested in the case where the correlation ρ

$$\rho = \frac{E\{[\Pi - M][V - \mu]\}}{S\sigma} \tag{10}$$

between the original portfolio and the new asset is small.

We sell a small amount λ of the portfolio and invest it in the new asset V:

$$\Pi' = (1 - \lambda)\Pi + \lambda V \tag{11}$$

A little algebra then shows that:

$$E\{\Pi'\} = (1-\lambda)M + \lambda\mu$$
$$SD\{\Pi'\} = \sqrt{(1-\lambda)^2 S^2 + 2(1-\lambda)\lambda\rho S\sigma + \lambda^2\sigma^2}$$
$$= (1-\lambda)S + O(\lambda^2) + O(\rho\lambda) \tag{12}$$

where $O(\lambda^2)$ denotes terms that are of the order of λ^2 and $O(\rho\lambda)$ denotes terms that are proportional to ρ times λ. Now if both ρ and λ are small, such as 1% or even 5%, and if σ is not huge compared with S, then these additional terms are negligibly small, and

$$SD\{\Pi'\} \oplus (1-\lambda)S \tag{13}$$

The Sharpe ratio of the new portfolio is

$$R_{Sharpe}\{\Pi'\} = \frac{(1-\lambda)M + \lambda\mu - r_f}{(1-\lambda)S} + O(\lambda^2) + O(\rho\lambda)$$
$$= \frac{M - r_f}{S} + \lambda\frac{\mu - r_f}{S} + O(\lambda^2) + O(\rho\lambda) \tag{14}$$

where r_f denotes the risk-free rate. The new Sharpe ratio is larger than that of the original one if and only if $\mu > r_f$.

We can generalize the above analysis to the case of new assets $V_1 \ldots V_n$ that are independent from (and thus uncorrelated with) both the original portfolio Π and among one another. We denote their means and standard deviations by:

$$E\{V_j\} = \mu_j$$
$$SD\{V_j\} = \sigma_j \tag{15}$$

We move amounts λ_j into the new assets V_j:

$$\Pi' = \left(1 - \sum_{j=1}^{n}\lambda_j\right)\Pi + \sum_{j=1}^{n}\lambda_j V_j \tag{16}$$

Then we have:

$$E\{\Pi'\} = \left(1 - \sum_{j=1}^{n}\lambda_j\right)M + \sum_{j=1}^{n}\lambda_j\mu_j$$
$$SD\{\Pi'\} = \sqrt{\left(1 - \sum_{j=1}^{n}\lambda_j\right)^2 S^2 + \sum_{j=1}^{n}\lambda_j^2\sigma_j^2}$$
$$= \left(1 - \sum_{j=1}^{n}\lambda_j\right)S + O(\lambda_j^2) \tag{17}$$

For the N-th central moment, we find:

$$M_{N-th}^{(central)}\{\Pi'\}$$
$$= \left(1 - \sum_{j=1}^{n}\lambda_j\right)^N M_{N-th}^{(central)}\{\Pi'\} + \sum_{j=1}^{n}\lambda_j^N M_{N-th}^{(central)}\{V_j\}$$
$$= \left(1 - \sum_{j=1}^{n}\lambda_j\right)^N M_{N-th}^{(central)}\{\Pi'\} + O(\lambda_j^N) \tag{18}$$

Therefore, the higher moments and thus the return probability distribution of the new portfolio are completely dominated by those of the remainder of the original portfolio.

To summarize our main results:

1. If we move only a very small amount into an independent (or very weakly correlated) new asset, we are mostly interested in the expected return of the new asset. The impact of its variance and its higher moments is highly suppressed.
2. If the new asset has an expected return larger than the risk-free rate, $\mu > r_f$, we can increase the Sharpe ratio of the portfolio by moving a small amount into the new asset.
3. To a very good approximation, the higher moments (in particular skewness and kurtosis) do not change. Therefore, if the original portfolio has a normal distribution, then the new portfolio is also approximately normal.

These results justify the use of the mean-variance approach when adding (small) CAT bond components to a diversified stock and/or bond portfolio, even though the distribution of the CAT bond returns is far from being normal.

Comparing Quoted Spreads for Hurricane-Linked Securities

Many CAT bonds pay interest according to the following conventions: Payments occur monthly, quarterly, or semiannually. For a given period, the cash flows are calculated as follows:

$$A/360 \times [\text{LIBOR} + s] \times \text{Principal} \tag{19}$$

The spread s is the spread over LIBOR as quoted in the offering circular. A is the actual number of days in the period, and the payment takes place at the end of the period. LIBOR is determined at the beginning of the period, with its maturity corresponding to the coupon frequency (one-month LIBOR if monthly payments, three-month LIBOR if quarterly, and six-month LIBOR if semiannually). The relevant data can be found in Exhibit 31.

Two issues should be considered when comparing these quoted spreads with other spreads. First, the "actual/360" convention used increases the cash flows, when compared with spreads quoted in "actual/actual." Second, many hurricane-related bonds have risk periods somewhat less than a full

EXHIBIT 31
Data Determining the Cash Flows for Selected Hurricane-Linked Securities

Transaction	Dated Date	Maturity	#d (total number of days over which interest accrues)	Tranche	Quoted Spreads	#d/360 x s	Coupon Payment Frequency
Res Re '97	06/16/97	06/15/98	364	Class A-1	273	276	Monthly
				Class A-2	576	582	
Trinity Re '98	03/03/98	12/31/98	303	Class A-1	186	157	Semiannually
				Class A-2	436	367	
Res Re '98	06/16/98	06/01/99	350	Notes	416	404	Quarterly
Mosaic I	07/17/98	07/09/99	357	Class A	444	440	Semiannually
				Class B	827	820	
				Principal-Protected	216.5	214.7	

year. Because of the seasonality of hurricanes, this does not reduce the actual risk (expected loss) to the note holders. Therefore, comparisons of risk numbers should take the number of days #d into account, over which interest will be paid. We should compare $\#d/360 \times s$, which is proportional to the actual cash flows.

The Law of Iterated Expectations

Divide the set of possible outcomes for a random variable X into two disjoint sets A and B:

$$A \cup B = \text{set of all possible outcomes}$$
$$p_A + p_B = 1 \qquad (20)$$

The law of iterated expectations states:

$$E[X] = p_A E_A[X] + p_B E_B[X] \qquad (21)$$

where E[X] is the unconditional expectation of X, and $E_A[X]$, $E_B[X]$ are the conditional expectations:

$$E_A[X] = E[X|A]$$
$$E_B[X] = E[X|B] \qquad (22)$$

By the law of iterated expectations, the variance of X is the sum of the conditional variances and the squared deviations of the conditional means from the unconditional mean, weighted by the probabilities p_A and p_B:

$$\begin{aligned} var[X] &= E[X^2] - (E[X])^2 \\ &= p_A var_A[X] + p_B var_B[X] \\ &+ p_A(E_A[X] - E[X])^2 + p_B(E_B[X] - E[X])^2 \end{aligned} \qquad (23)$$

where

$$var_A[X] = E_A[X]^2 - p_B(E_B[X] - E[X])^2 \qquad (24)$$

Proof:

$$\begin{aligned} & p_A var_A[X] + p_B var_B[X] + \\ & \quad p_A(E_A[X] - E[X])^2 + p_B(E_B[X] - E[X])^2 \\ &= p_A\{E_A[X^2] - (E_A[X])^2 + (E_A[X])^2 - \\ & \quad 2E_A[X]E[X] + (E[X])^2\} + \\ & \quad p_B\{E_B[X^2] - (E_B[X])^2 + (E_B[X])^2 - \\ & \quad 2E_B[X]E[X] + (E[X])^2\} \\ &= p_A E_A[X^2] + p_B E_B[X^2] - \\ & \quad 2(p_A E_B[C] + p_B E_B[X])E[X] + \\ & \quad (p_A + p_B)(E[X])^2 \\ &= E[X^2] - 2(E[X])^2 + (E[X])^2 \\ &= E[X^2] - (E[X])^2 \\ &= var[X] \end{aligned} \qquad (25)$$

ENDNOTES

The authors especially thank Robert Litzenberger, now head of firmwide risk management at Goldman, Sachs & Co., who initiated this article and provided important suggestions and comments. His knowledge of insurance-related products and financial economics has been a source of inspiration and guidance for the research in this field at Goldman Sachs. We also received extensive and very important comments on the manuscript from members of the Risk Markets Group (formerly the Insurance Products Group) at Goldman Sachs: Vivek Bantwal, Gaurav Bhandari, Simon Collier, Vijay Culas, Daniela

Gatto, Mark Griffin, Per Johansson, Andrew Kaiser, Shiv Kumar, Bryce Markus, Michael Millette, Barney Schauble, and Morgan Sze. The authors are grateful to Ronald A. Krieger and William Tompkins for their help in editing the manuscript. Last but not least, they also thank Han Chen, Weimin Dong, Mark Romer, Rick Wales, and Chesley Williams at Risk Management Solutions for their help in preparing and reviewing this report.

[1]It is possible for a company to issue insurance-linked securities directly or to engage in similar derivative transactions, by foregoing certain advantages associated with use of an SPV (such as achieving reinsurance accounting treatment). In addition, though most offerings to date have supplied risk capital to insurance or reinsurance companies, this SPV structure has been used to provide insurance coverage to corporate entities as well.

[2]The adjusted historical approach was the basis for Pielke and Landsea [1998]. It is interesting to compare their results with those obtained by expert modeling firms through ex-post simulations, although it is first necessary to adjust Pielke and Landsea's figures to insured rather than total economic losses.

[3]According to the "HURDAT" catalog, during the period 1899–1997, 157 hurricanes made landfall in the continental United States. The HURDAT catalog is maintained by the National Hurricane Center/Tropical Prediction Center (NHC/TPC) and is updated annually.

[4]In particular, see Gray et al. [1998] in their forecasts of Atlantic seasonal hurricane activity.

[5]See, for example, the Open File Report 88-398 of the Working Group on California Earthquake Probabilities.

[6]All of the industry loss numbers in this section are sourced from Table 6, page 36, of the Swiss Re, sigma No. 4/1998 report, and are in 1997 dollars.

[7]The spreads of high-yield bonds displayed in Exhibit 12 correspond to the summer of 1998. Note that the relevant comparison is between high-yield bonds and CAT bonds at the time of issuance of the respective CAT bonds. Using the same method as described in the text, we estimate current (summer 1999) numbers as follows: Spreads over LIBOR are 123, 183, 342, 370, and 486 bp for bonds rated Ba2, Ba3, B1, B2, and B3, respectively (assuming a swap spread of 50 bp). This implies Sharpe ratios of 0.30, 0.08, 0.14, 0.02, and –0.11.

[8]The formation of a hurricane in the Atlantic can be observed over weeks, but crucial information, including where it makes landfall and the event characteristics, will be available only shortly before it actually hits the coast.

[9]Arguably, the events of the last year have shown that prices of emerging market bonds can also exhibit large jumps.

[10]Certain other financial risks — notably temperature — will exhibit a more typical, gradual price process, which should be attractive to investors concerned about jump risk.

[11]Skewness, kurtosis, and "normalized" higher moments are defined as dimensionless numbers by dividing the central moments by the suitable power of the standard deviation. Skewness is the third central moment divided by the cube of the standard deviation, kurtosis is the fourth central moment divided by the fourth power of the standard deviation, and so on.

REFERENCES

Gordon, Richard T., and Auguste Boissonnade. "The Return of El Niño." *Solutions*, Vol. 1, Number 1 (Autumn 1997).

Gordon, Richard T., Auguste Boissonnade, and J.B. Elsner. "Track Classification of North Atlantic Tropical Cyclones." Submitted to *Monthly Weather Review*, April 1998.

Gray, William M., et al. "Forecast of Atlantic Seasonal Hurricane Activity for 1998," and previous forecasts for 1994-1997. Published on the Internet at http://typhoon.atmos.colostate.edu/forecasts.

Litzenberger, Robert H., David R. Beaglehole, and Craig E. Reynolds. "Assessing Catastrophe-Reinsurance-Linked Securities as a New Asset Class." *Journal of Portfolio Management*, Special Issue (Fisher Black Tribute), December 1996.

Moody's Investors Service Inc. *Historical Default Rates of Corporate Bond Issuers, 1920–1997* (1998).

Pielke, Roger A., and Christopher W. Landsea. "Normalized Hurricane Damages in the United States: 1925–1995." *Weather and Forecasting*, September 1998.

Swiss Re. "Too Little Reinsurance of Natural Disasters in Many Markets." *sigma*, No. 7/1997.

Swiss Re. "Natural Catastrophes and Major Losses in 1997: Exceptionally Few High Losses." *sigma*, No. 4/1998.

Working Group on California Earthquake Probabilities. "Probabilities of Large Earthquakes Occurring in California on the San Andreas Fault." Open File Report 88-398, U.S. Department of the Interior, 1988.