THE BASIS RISK OF CATASTROPHIC-LOSS INDEX SECURITIES

By

J. David Cummins, David Lalonde, Richard D. Phillips

January 11, 2002

Please address correspondence to: Richard D. Phillips
Department of Risk Management and Insurance
Georgia State University
P.O. Box 4036
Atlanta, GA 30302-4036

J. David Cummins
The Wharton School
Phone: 215-898-5644
Fax: 215-898-0831
cummins@wharton.upenn.edu

David Lalonde
Applied Insurance Research
Phone: 425-990-4703
Fax: 206-583-8382
dlalonde@air-boston.com

Richard D. Phillips
Georgia State University
Phone: 404-651-3397
Fax: 404-651-4219
rphillips@gsu.edu

THE BASIS RISK OF CATASTROPHIC-LOSS INDEX SECURITIES

ABSTRACT

This paper provides an analysis of the basis risk of catastrophic-loss (CAT) index options, which securitize losses from events such as hurricanes and earthquakes. These securities were introduced in response to a surge in the frequency and severity of catastrophic losses and the recognition that global reinsurance markets are not efficient in financing this type of loss. This paper analyzes the effectiveness of CAT index options in hedging hurricane losses for 255 insurers accounting for 93 percent of the insured residential property values in Florida. We simulate hurricane losses using a sophisticated model developed by Applied Insurance Research and measure the basis risk of call spread hedges that settle on a statewide loss index and four intra-state indices. The results show that insurers in the two largest size quartiles can hedge losses almost as effectively using the intra-state index contracts as they can using contracts that settle on their own losses. Many insurers in the third largest size quartile also can hedge effectively using the intra-state indices, but most insurers in the smallest quartile would encounter significant basis risk. A high proportion of the total property exposure in Florida can be hedged efficiently using the intra-state contracts, providing a potential solution to the CAT loss financing problem. Hedging using the statewide loss index is effective only for the largest and most diversified insurers.

The Basis Risk of Catastrophic-Loss Index Securities

1. Introduction

An important recent innovation in financial markets is the securitization of losses from catastrophic (CAT) events such as hurricanes and earthquakes. The development of these instruments has been motivated by a surge in the frequency and severity of catastrophic losses. Hurricane Andrew in 1992 and the Northridge earthquake in 1994 resulted in $30 billion in insured property losses, and recent projections indicate that the losses from a major Florida hurricane or California earthquake could exceed $100 billion.[1] Losses of this magnitude would significantly stress the capacity of the insurance industry, but are small relative to the size of U.S. stock and bond markets.[2] Thus, securitization offers a potentially more efficient mechanism for financing CAT losses than conventional insurance and reinsurance (Jaffee and Russell 1997, Froot 1998a, 2001). Both insurers and non-insurers such as industrial firms can use these instruments to hedge their exposure to catastrophic losses, in effect permitting the non-insurers to bypass the insurance market.[3] Moreover, because catastrophic losses are "zero-beta" events, CAT-loss securities provide a valuable new source of diversification for investors (Litzenberger, et al. 1996, Canter, et al. 1997).

CAT-risk securities offer a particularly interesting example of a new type of derivative where the underlying is not a traded asset or commodity, so that prices are not observed.[4] In the absence of a traded underlying asset, insurance-linked securities have been structured to pay-off on three types of variables – insurer-specific catastrophe losses, insurance-industry catastrophe loss indices, and *parametric* indices based on the physical characteristics of catastrophic events. An important consideration in the choice between an

---

[1]Unpublished data from Applied Insurance Research, Boston.

[2]A loss of $100 billion would equal approximately 30 percent of the equity capital of the U.S. property-liability insurance industry but would be less than 0.5 of 1 percent of the value of U.S. stock and bond markets.

[3]CAT securities also enable insurers and non-financial firms exposed to CAT risk to hedge losses exceeding the capacity of the international insurance and reinsurance markets and to avoid the market disruptions caused by reinsurance price and availability cycles (Cummins and Weiss 2000, Froot 2001).

[4]In this regard, CAT securities are analogous to other new derivatives with "exotic underlyings," such as weather derivatives (Geman 1999).

insurer-specific trigger and an industry loss-index or parametric trigger is the relative cost of moral hazard versus basis risk (Doherty 1997). Securities based on insurer-specific (or hedger-specific) losses have no basis risk but expose investors to moral hazard; whereas securities based on industry loss indices or parametric triggers greatly reduce or eliminate moral hazard but expose hedgers to basis risk. Because industry loss-index contracts are also expected to have lower transactions costs and higher liquidity than insurer-specific contracts,[5] index-linked contracts could come to dominate insurer-specific contracts for many insurers, provided that basis risk is sufficiently low.

CAT risk securitization began in 1992 with the introduction of index-linked catastrophic loss futures contracts by the Chicago Board of Trade (CBOT). The CBOT contracts evolved into call option spreads, but were later withdrawn because of low trading volume. The majority of risk capital raised to date has been generated through the issuance of CAT bonds that settle on the losses of the issuing insurer.[6]

A primary reason for the lack of success of the CBOT index contracts and the general predominance of insurer-specific contracts so far is the perception among insurers that index-linked CAT securities are subject to unacceptable levels of basis risk (American Academy of Actuaries 1999).[7] However, there is virtually no empirical evidence about the basis risk of index-linked securities when used in hedging

[5]Loss-index contracts have an advantage in terms of transactions costs and liquidity because it is easier to standardize contracts and report losses on an index, versus potentially having a range of contract specifications and triggering criteria depending upon the characteristics of the issuer as with insurer-specific contracts.

[6]The proceeds of a CAT bond issue are held by a single purpose reinsurer and invested in safe securities such as Treasury bonds. If a specified catastrophic event occurs, the hedger can use the bond proceeds to offset catastrophic losses; and there is full or partial forgiveness of the repayment of principal and/or interest. If the defined catastrophic event does not occur, the investors receive their principal plus interest equal to the risk free rate plus a risk-premium. Among the earliest successful CAT bonds were those issued by Winterthur Re, USAA, and Swiss Re in 1997. The first successful CAT bond issued by a non-financial firm, occurring in 1999, covers Tokyo earthquake losses for Oriental Land Company, Ltd. For further discussion see Froot (2001).

[7]CBOT contracts were available for nine indices – a national index, five regional indices, and three state indices (for California, Florida, and Texas). Even the single-state indices were viewed by many insurers as potentially having excessive basis risk. In this study, we look both at Florida loss indices and at four sub-state indices to provide evidence on whether insurer concern about basis risk of state-wide contracts is justified and on the degree to which sub-state indices can be used to reduce basis risk and improve hedging effectiveness.

catastrophe risk by specific insurers. The primary objective of this paper is to remedy this deficiency in the existing literature by conducting a comprehensive analysis of the basis risk and hedging-effectiveness of index-linked CAT loss securities. We conduct a simulation analysis of hedging-effectiveness for 255 insurers accounting for 93 percent of the insured property values in Florida, the state with the highest exposure to hurricane losses. The study is based on data provided by the Florida Insurance Commissioner on county-level insured residential property values for each insurer in the sample.

The study proceeds by simulating hurricane losses for each insurer in the sample using a sophisticated model developed by Applied Insurance Research (AIR), a leading CAT modeling firm.[8] The AIR hurricane model combines actuarial data, vulnerability relationships for various construction types, historical climatological data, and meteorological models of the underlying physical processes that drive the severity and trajectory of hurricanes. We use the AIR model to obtain estimates of insurer losses over a simulation period consisting of 10,000 years of hurricane experience. We then utilize the simulated loss experience to analyze the effectiveness of catastrophic loss hedging strategies for the sample insurers.

The analysis focuses on non-linear hedging strategies where the hedge portfolio consists of a short position in catastrophe losses and a long position in call option spreads on a CAT loss index. We analyze non-linear hedging because the call-spread is the functional form for payoffs on nearly all CAT bonds and options issued to date as well as for conventional catastrophe reinsurance contracts. Several hedging objectives are investigated, including reduction in loss volatility (variance), value-at-risk (VaR), and the expected loss conditional on losses exceeding a specified loss threshold. Consistent with theory (e.g., Raviv 1979), most of the analysis is conducted under the assumption insurers seek to hedge relatively large losses, based on the rationale that hedging is costly so that insurers prefer to retain loss volatility due to smaller

---

[8]The AIR model has been widely used by insurers and reinsurers since 1987 in monitoring their exposure to catastrophic losses and developing underwriting strategies and was the first model to meet the standards of the Florida Insurance Commission on Hurricane Loss Projection Methodology.

events or to use risk management strategies that are less expensive than purchased hedging instruments.[9]

The benchmark model of hedging effectiveness is the *perfect hedge*, defined as the risk reduction a hedger could achieve for large loss events by using its own loss experience as the hedge index. The perfect hedge is roughly equivalent to purchasing catastrophe reinsurance or issuing insurer-specific CAT bonds.[10] The effectiveness of the perfect hedge is compared with hedges based on a statewide loss index, analogous to the CBOT's Florida index, and four intra-state regional indices. The analysis measures the degree of basis risk insurers would incur from hedging through index-linked CAT loss securities.

By way of preview, the principal finding of our study is that insurers in the two largest size quartiles can hedge large losses effectively using intra-state regional indices.[11] Many insurers in the third largest size quartile also can hedge effectively using the intra-state indices, but hedging by insurers in the smallest size quartile is significantly less effective. Finally, although the results show that many insurers would encounter significant basis risk in hedging with a state-level index, even with this index a high proportion of the total property value exposed to loss in Florida could be hedged efficiently; and an even higher proportion of exposures could be hedged efficiently using regional indices.

The findings are important as a case study in the securitization of a non-traded asset, and thus can provide guidance for the securitization of other unconventional financial exposures. Our methodological approach also has the potential to serve as a model for analyzing the hedging effectiveness of other securities on exotic underlyings, such as weather derivatives. The results have important implications for insurers, not only with respect to hedge efficiency but also for the management of underwriting exposure. The analysis

[9]Froot (2001) provides evidence that both reinsurance and CAT bonds generally trade at significant margins above the expected loss covered by the hedge.

[10]Reinsurance and CAT bond contracts are not in fact "perfect" because such instruments usually include provisions such as coinsurance requirements to control moral hazard (see further discussion below).

[11]Throughout the paper, size quartiles are based on total residential property value exposed to loss in Florida. Quartile 1 is the largest size quartile and quartile 4 the smallest.

should be of interest to insurance regulators and policy makers concerned about financing losses from catastrophic events and preventing the destabilization of insurance markets due to catastrophes.

There have been two previous empirical studies of the basis risk of insurance-linked securities, both using different or less comprehensive study designs. Harrington and Niehaus (1999) conduct a time series analysis of the correlation between state-specific loss ratios for a sample of insurers and the CAT loss index compiled by Property Claims Services (PCS), an insurance industry statistical agent, and find that PCS-index derivatives would have provided effective hedges for many homeowners insurers. In a study more similar to ours, Major (1999) conducts a simulation analysis of CAT losses based on insurer exposures in Florida and finds that hedging with a statewide CAT index is subject to substantial basis risk. Our analysis extends Major's by considering much larger numbers of insurers and storms, testing intra-state indices as well as a statewide index, and evaluating a wider variety of hedging strategies.

The remainder of the paper is organized as follows: Section 2 describes the data, the AIR model, and the basis risk study design. The results of the analysis are presented in section 3, and section 4 concludes.

2. Data and Basis Risk Study Design

The basis risk study has three major phases: (1) The identification and analysis of data on the catastrophic loss exposure of a sample of insurance companies. (2) The simulation of catastrophic losses in the geographical area covered by the sample companies. (3) The measurement of basis risk and hedge effectiveness for the insurers in the sample using a variety of hedging strategies and loss indices.

The Data

The data base for the study consists of county-level data, obtained from the Florida Insurance Commissioner, on insured residential property values for 255 of the 264 insurers writing property coverage in Florida in 1998.[12] The insurers in our sample account for 93 percent of the total insured residential property values in the state. Thus, our results can be interpreted as generally representative of the entire

[12]Data on the nine omitted insurers were not available from the Florida Insurance Commissioner.

insurance industry. Further details about the sample are provided in the section 3.

Catastrophic Loss Simulations

The simulated catastrophic losses for our sample of insurers are generated using the hurricane model developed by Applied Insurance Research. This section provides a brief description of the model. Further details on the model are provided in Appendix A and in Applied Insurance Research (1999).

The hurricane loss estimation methodology employed by AIR is based on well-established scientific theory in meteorology and wind engineering. The simulation models were developed through careful analyses and synthesis of all available historical information and incorporate statistical descriptions of a large number of variables that define both the originating event (e.g., hurricane) and its effect on insured structures. The models are validated and calibrated through extensive processes of both internal and external peer review, post-disaster field surveys, detailed client data from actual events, and overall reasonability and convergence testing. The AIR hurricane model has been used by the insurance industry since 1987 and is well known for its reliability and the credibility of the loss estimates it generates. The AIR model was the first to meet the standards of the Florida Insurance Commission on Hurricane Loss Projection Methodology.

The structure of the simulation model is summarized in Figure 1. The process begins with a Monte Carlo simulation of the number of storms per year for a 10,000 year simulation period, generating more than 18,000 simulated events. The landfall and meteorological characteristics are then simulated for each storm, where the meteorological characteristics include central barometric pressure, radius of maximum winds, forward speed, storm direction, and storm track. Once the model generates the storm characteristics and point of landfall, it propagates the simulated storm along a path characterized by the track direction and forward speed. To estimate the property losses resulting from the simulated storms, the AIR hurricane model generates the complete time profile of wind speeds, or windfield, at each location affected by the storm.

After the model estimates peak wind speeds and the time profile of wind speeds for each location, it generates damage estimates for different types of property exposures by combining data on insured property values and structure characteristics with wind speed information at each location affected by the event. To estimate building damage and the associated losses, the AIR hurricane model uses damageability relationships, or damage functions which have been developed by AIR engineers for a large number of building construction and occupancy classes. In the last component of the catastrophe model, insured losses are calculated by applying the policy conditions to the total damage estimates. Policy conditions include deductibles, coverage limits, coinsurance provisions, and a number of other factors.

A fundamental component of the model is AIR's insured property data base. AIR has developed databases of estimated numbers, types, and insured values for residential, commercial, mobile home, and automobile properties in the United States by five-digit ZIP code. These databases have been constructed from a wide range of data sources and reflect the estimated total replacement cost of U.S. property exposures. In the present study, AIR's zip code level data on insured property values for companies doing business in Florida were used in the simulations and aggregated to the county level using information supplied by the Florida Insurance Department to protect the confidentiality of AIR's data bases. The simulations were also conducted using the AIR zip-code data base exclusively for a random sample of five companies in order to validate the county aggregation approach. The validation tests indicated that aggregating our results to the county level provides an accurate representation of the losses that would have been generated using AIR's zip code data base as the exclusive source of information.

Hedging Strategies and Hedge Effectiveness

The objective of the hedging analysis is to determine the effectiveness of hedges based on a statewide loss index and four intra-state regional indices. The four intra-state indices are based on a subdivision of the state into four segments – the Panhandle, Gulf Coast, North Atlantic, and South Atlantic (see Appendix B). The segment definitions were provided by Applied Insurance Research based on experience with insurance clients including analyses conducted in conjunction with the United Services Automobile Insurance (USAA) CAT bond issues of 1997-1999. Four regions were chosen as a subdivision of the state that we hypothesized would be sufficient to enable insurers to create effective hedges without incurring the high transactions costs and lack of liquidity that would likely result from a finer subdivision of the state.[13] The subdivision of the state was not optimized to minimize basis risk. Therefore, it is possible that a geographical subdivision could be found that would provide more effective hedges than the four regions used in our analysis.

Index-hedge effectiveness is measured relative to the performance of *perfect hedges*, which pay off on the insurer's own losses. The perfect hedge parallels the results the insurer could attain by purchasing conventional reinsurance contracts or issuing insurer-specific CAT bonds, whereas the index hedges are designed to reflect results that could be achieved through trading in index-linked CAT options.

The analysis assumes that insurers are risk-neutral but are motivated to hedge by market imperfections, including direct and indirect costs of financial distress and convex tax schedules.[14] In addition, because the role of insurance is to indemnify policyholders for insured losses, insurers are motivated to maintain a reputation for having low default risk. In this regard, risk management can be viewed as a substitute for holding costly equity capital.

We expect insurers to purchase securitized hedging instruments to perform a similar function as traditional excess-of-loss (XOL) catastrophe reinsurance, i.e., hedging of losses in the tail of the loss severity distribution that are most likely to disrupt operations and threaten solvency. Thus, we primarily analyze *large loss* or *conditional hedging,* where hedges are constructed to reduce risk conditional on the loss exceeding a given amount or percentile of the loss distribution. Large loss hedging is consistent with observed behavior in the CAT bond market and is also consistent with corporate risk management theory, which postulates that firms are motivated to hedge to avoid financial distress costs and to maintain internal

[13]A 1998 attempt to launch zip-code level index contracts failed to generate interest among insurers and investors and is currently dormant. Chookaszian and Ward (1998) discuss the proposed indices.

[14]For more extensive discussions of the rationale for corporate risk management, see Merton and Perold (1993) and Froot, Scharfstein, and Stein (1993). Because insurers purchase reinsurance at prices that are usually significantly higher than expected costs (Froot 2001), they demonstrate a revealed preference for hedging activities that would generally not make sense with frictionless markets and risk neutrality.

capital to finance future growth opportunities. The analysis thus assumes that insurers are able either to retain the risk from relatively small loss events or to manage this risk using less expensive alternatives than reinsurance or securitized hedging instruments.[15]

We consider "buy and hold" hedging strategies covering a single period, because this is the standard approach used by insurers when purchasing excess of loss reinsurance contracts and issuing CAT bonds. We analyze non-linear hedges, where the insurer forms a hedge portfolio consisting a short position in unhedged catastrophe losses and a long position in call option spreads. The non-linear analysis is emphasized because the call option spread is the dominant contractual form in both the catastrophe reinsurance and CAT securities markets (see Froot 1998b, 2001, Cummins, Lewis, and Phillips 1999).[16]

Call-Spread Hedging

As discussed above, the insurer is assumed to hedge by forming a portfolio consisting of its own unhedged catastrophic losses and a position in call option spreads on a loss index. Defining insurer j's unhedged loss as L_j and its hedged net loss under loss index i as L_j^i, insurer j's loss under the perfect hedge (i = P) is:

$$L_j^P = L_j - h_j^P \left[Max(L_j - M_j^P, 0) - Max(L_j - U_j^P, 0) \right] \tag{1}$$

[15]If natural disasters are zero-beta events and significant market imperfections are not present, the rate of return on CAT bonds should approximately equal the risk-free rate plus a risk premium sufficient to compensate investors for the expected loss of principal due to a triggering catastrophic event. The CAT bonds issued to date, however, have been priced at spreads equal to several times the expected loss (the median risk-premium to expected-loss ratio is about 6.8, see Appendix Table A.1). Possible explanations for the high spreads include moral-hazard, the illiquidity of the bonds, uncertainty about expected loss estimates, and investor unfamiliarity with the contracts. For further discussion, see Kunreuther and Bantwal (1999).

[16]For purposes of comparison with prior work (e.g., Harrington and Niehaus 1999, Major 1999), we also analyze linear hedging strategies where hedge portfolios are formed that linearly combine a short position in CAT losses with a long position in CAT loss futures. This is the type of portfolio evaluated in most of the existing hedging literature for other types of risks (e.g., Ederington 1979). The results of this analysis, which are available from the authors, lead to similar conclusions about the hedging effectiveness of Florida loss indices but are less relevant here than the non-linear results because linear hedges are not used in the CAT risk market.

where h_j^P = the hedge ratio for the perfect hedge, M_j^P = the lower strike price of the call spread, and U_j^P = the upper strike price of the spread.

The perfect hedge is compared to hedges based on loss indices that are not perfectly correlated with the insurer's losses. Insurer j's net loss based on an index consisting of industry-wide, state-level losses is:

$$L_j^S = L_j - h_j^S [Max(L^S - M_j^S, 0) - Max(L^S - U_j^S, 0)] \tag{2}$$

where L_j^S = insurer j's hedged loss using an state-wide industry loss index, h_j^S = the hedge ratio for the state-wide hedge, $L^S = \Sigma_j L_j$ = state-wide losses for the industry, and M_j^S and U_j^S are the lower and upper strike prices for company j's state-wide call spread. Insurer j's hedged loss under the intra-state hedge is:

$$L_j^R = \sum_{r=1}^{R} \left[L_{jr} - h_j^r [MAX(L_r^R - M_j^r, 0) - MAX(L_r^R - U_j^r, 0)] \right] \tag{3}$$

where L_j^R = company j's losses under the intra-state regional hedge, L_{jr} = the unhedged losses of insurer j in region r, h_j^r = hedge ratio for insurer j in region r, L_r^R = industry-wide losses in region r, M_j^r = lower strike price for company j's region r call option spread, and U_j^r = upper strike price for company j's region r call spread, and R = the number of regions = 4 in our analysis).

Interviews with executives of insurance and reinsurance companies, conducted by the authors in the course of this research project, revealed that insurers are concerned about collecting less on an index hedge than they would under a reinsurance contract for a catastrophic event. Accordingly, we also analyze the *index hedge basis*, defined as follows for the regional hedge:

$$B_j^R = (L_j^P - L_j^R) \tag{4}$$

where B_j^R = the regional-hedge basis for insurer j. The state-hedge basis is defined similarly. Equivalently, the basis is equal to the difference between the payment on the index hedge and the payment on the perfect hedge. Therefore, negative values for the basis imply under-collection relative to the perfect hedge and positive values imply over-collection (basis gain).

In the general call-spread hedging problem, the insurer is assumed to minimize a function of L_j^i subject to a cost constraint, conditional on state-wide losses exceeding a specified large-loss threshold. Defining the objective function for criterion m and index i as $G_m(L_j^i)$, the optimization problem using a state-wide hedge is given by:

$$\underset{h_j^S,\, M_j^S,\, U_j^S}{Minimize} \quad G_m[L_j^S \,|\, L_j \in \{L_j \,|\, L^S > T\}] \tag{5}$$

Subject to: $h_j^S\,[W(L^S,M_j^S) - W(L^S,U_j^S)] \leq C_j$

where C_j = the maximum amount available to insurer j to spend on hedging, and $W(L^S,M_j^S)$ and $W(L^S,U_j^S)$ = the prices of call options on industry losses L^S with strike prices M_j^S and U_j^S, respectively. Thus, the insurer optimizes by choosing the hedge ratio and the two strike prices, M_j^S and U_j^S, subject to spending a maximum of C_j on hedging. By varying C_j, it is possible to generate an efficient frontier based on each optimization criterion and loss index. The optimization problem for the perfect hedge is defined similarly to expression (5). The optimization problem for the regional hedge is analogous to (5) except that there are twelve decision variables – four hedge ratios and four sets of lower and upper strike prices.

In the optimization problems involving the basis, B_j^i, the basis is substituted for L_i^j in (5), and an additional conditioning constraint is imposed. That is, the optimization is conditional on $L_j^P \neq L_j$, so that the optimization criterion function takes into account only those cases where payment is triggered under the perfect hedge. It is this set of cases that is relevant in evaluating basis risk in terms of B_j^i.

In keeping with our large-loss hedging strategy, most of the optimization problems are solved over the subset of losses that are generated by an <u>industry-wide</u> loss event such that state-wide losses ($L^S = \sum_i L_i$) exceed a specified threshold (T). Using the industry-wide conditioning criterion enables us to standardize the comparisons among hedges by solving the hedging problem over the same set of losses, i.e., losses enter into the optimization problems by virtue of being generated by an industry-wide loss at the state level that exceeds the threshold.[17] The optimization problems are solved for thresholds of T = \$1 billion, \$2.5 billion, and \$5 billion and also for the unconditional case, where T = 0. The three non-zero thresholds correspond, respectively, to the 23rd, 14th , and 8th percentiles of the Florida residential property loss severity distribution.

We focus on three hedging criteria which are either standard in the hedging literature or likely to be appropriate for insurers: (1) the variance of net losses or the basis (conditional on losses being in the set of losses that exceed the threshold), (2) the value-at-risk (VaR), and (3) the expected exceedence value (EEV). Conditional variance reduction is the most straightforward of the three hedging criteria, giving rise to the objective function: $G_{1T}(L_j^i * L_j \, 0 \, L^T) = \sigma_T^2[L_j^i(h_j^i,M_j^i,U_j^i) * L_j \, 0 \, L^T]$ = the conditional variance of insurer j's loss net of the payoff on the call option spread using loss index i and threshold T, where $L_j^i(h_j^i,M_j^i,U_j^i)$ = insurer j's hedged loss using hedge index i with hedge ratio h_j^i and strike prices M_j^i and U_j^i. The hedge index i = P for the perfect hedge, S for the statewide industry hedge, and R for the intra-state regional hedge, where the latter is a function of twelve rather than three variables. L^T is used as an abbreviation for $(L_j * L^S > T)$. The conditional variance of the basis is defined similarly, with the additional event constraint, $L_j^P \ldots L_j$

Value-at-risk (VaR) reduction has received considerable attention in the hedging and financial risk management literature (e.g., Ahn, et al. 1999, Dowd 1999, Santomero 1997).[18] VaR is defined as the amount of loss such that the probability of exceeding this amount during a specified period of time is equal to α, a small positive number (0 < α < 1). Stated more formally, defining insurer j's net loss distribution function under hedge index i and threshold T as $F_{ijT}[L_j^i(h_j^i,M_j^i,U_j^i) * L_j \, 0 \, L^T] = F_{ijT}(\cdot)$, VaR is defined as:

---

[17]The exception to this general rule is the expected exceedence value (EEV) hedge, which is based on losses exceeding specified percentiles of the individual-insurer loss distributions. EEV-hedging was treated this way because we considered use of the individual insurer's loss distribution to be more consistent with the EEV concept, which is discussed in more detail below. As a robustness check, we also conducted the analysis for the other hedging criteria (variance and value at risk) under the assumption that insurers form hedges based on upper percentiles of their own loss distributions rather than the set of losses resulting from large industry-wide events and also conducted the EEV analysis using the industry loss distribution criterion. Use of these alternative hedging approaches had no material effect on the results.

[18]Moreover, VaR is similar in concept to the probability of ruin, which has been studied extensively by actuaries. Hence, insurers are likely to find VaR to be a familiar and informative criterion.

$$VaR_{ijT}[\alpha, L_j^i(h_j^i, M_j^i, U_j^i) | L_j \in L^T] = F_{ijT}^{-1}(1 - \alpha) \tag{6}$$

where $F_{ijT}^{-1}(\cdot)$ = the inverse of the net loss distribution function. Using VaR, the optimization function in expression (5) becomes $G_2(L_j^i | L_j \in L^T) = VaR[\alpha, L_j^i(h_j^i, M_j^i, U_j^i) | L_j \in L^T]$.

Although VaR is an important and useful statistic, in many cases the risk manager would like to know not only the probability that a given loss level will be exceeded but also the expected amount of loss conditional on the loss level being exceeded. This is the quantity measured by our third optimization criterion, the expected exceedence value (EEV). EEV is essentially the value of a call option on L_j^i with strike price equal to a specified loss threshold.[19] More formally, the EEV is defined as:

$$EEV_j[L_j^i(h_j^i, M_j^i, U_j^i) | L_j > L_V] = \int_{L_V}^{\infty} [L_j^i - L_V]\, dF_j(L_j^i(h_j^i, M_j^i, U_j^i)) \tag{7}$$

where L_V = a loss threshold specified by the decision maker. The EEV criterion function is $G_3(L_j^i) = EEV_j[L_j^i(h_j^i, M_j^i, U_j^i) | L_j > L_V]$. Thus, the insurer minimizes the expected excess loss conditional on the loss being equal to or greater than a specified loss threshold. This measure is more informative than the VaR in the sense that the risk manager is likely to care whether the threshold loss level is exceeded by \$1 or \$1 million. The VaR and EEV for the index-hedge basis are defined similarly.

Note that conditioning is implicit in the EEV, so it satisfies our criterion for large-loss hedging as long as L_V is sufficiently high. In our analysis, L_V was set at the 92.5th, 95th and 97.5th percentiles of each insurer's loss distribution. Conditioning on the insurers' own loss distributions in the EEV calculations seems more appropriate than using the industry loss distribution because the EEV uses the entire loss distribution above the attachment point.

For each loss index i, we define <u>hedge effectiveness</u> as the proportionate reduction in the unhedged

---

[19] Recent research suggests that EEV-type measures have desirable properties not possessed by value at risk measures. See, for example, Artzner, et al. (1999).

value of the criterion function. We denote the hedge effectiveness measure for insurer j based on loss index i as HE_{jm}^{i}, where m = 1, 2, and 3 for the variance, VaR, and EEV criteria, respectively. Under the EEV criterion function, for example, the hedge effectiveness of the state-wide index is:

$$HE_{j3}^{S} = 1 - \frac{EEV_j[L_j^S(h_j^S, M_j^S, U_j^S)|L_j > L^T]}{EEV_j[L_j|L_j > L^T]} \tag{8}$$

Hedge effectiveness is defined similarly for the other two hedging objectives and for the basis. Another useful indicator of hedge performance is hedge efficiency, defined as the hedging effectiveness of the index hedge relative to that of the perfect hedge, i.e.,

$$RHE_{jm}^{i} = \frac{HE_{jm}^{i}}{HE_{jm}^{P}} \tag{9}$$

where i = S = statewide hedge and i = R = regional hedge. Thus, whereas hedge effectiveness provides an absolute measure of hedge performance, hedge efficiency measures hedge performance relative to that of the perfect hedge and thus provides a better measure of the degree of basis risk than hedge effectiveness. Analysis of the hedge basis (equation (4)) provides an alternative measure of hedge efficiency.

Estimation Methodology

In solving the optimization problems discussed above, we adopted an estimation strategy that includes the use of both a standard calculus-based algorithm and a *differential evolutionary genetic algorithm* (Goldberg 1989).[20] The evolutionary genetic algorithm is a global optimization technique designed to efficiently investigate the "entire" feasible set of the parameter space.[21] This differs from the objective of conventional calculus-based optimizers, which seek to refine an initial "guess" about the optimal solution

---

[20]See Kingdon and Feldman (1995), Varetto (1998), and Engle and Manganelli (1999) for other applications of genetic algorithms in solving financial problems.

[21]Most global algorithms rely on random sampling and other techniques to efficiently search the feasible region rather than performing a deterministic uniform grid search. The latter approach tends to be computationally demanding when the parameter space is multi-dimensional.

vector. Whereas conventional optimizers tend to find the optimal solution in the region of the parameter space where the algorithm is started, global optimizers tend to range across alternative regions of the parameter space seeking the "region of attraction" that contains the global optimum. Global optimizers are less likely than conventional algorithms to converge to local optima and have been shown to possess global convergence properties under mild assumptions (Pinter 1996).

Although global optimization algorithms such as the differential genetic algorithm are superior to conventional methods in extensively exploring the space of possible solutions, they are not necessarily as efficient as conventional methods at refining the solution once the region where the global optimum is located has been identified (Pinter 1996, Goldberg 1989). Consequently, we adopt a two-stage estimation strategy whereby we first find a solution using the genetic algorithm and then use the results as starting values for a calculus-based optimization using the Newton-Raphson method.[22] Once Newton-Raphson results have been obtained for all cost constraints for a company/objective function combination, we conduct a visual inspection to check for discontinuities in the firm's variance, VaR, or EEV reduction frontier. Where discontinuities are found, the optimization process is rerun for the discontinuous case(s) to obtain results that are consistent with those for the other cost constraints in the risk-reduction frontier.[23]

---

[22]Combining global optimization and more conventional algorithms is referred to as a *hybrid optimization algorithm* in the optimization literature (Pinter 1996). Our hybrid methodology can be described in more detail as a five step approach: (1) Beginning with the lowest cost constraint, we select initial starting values for the lower and upper strike prices equal to the 97.5th and 99.5th percentiles of the insurer's loss distribution in the state or regions. These starting values were chosen based on preliminary experiments and on theoretical predictions that optimal hedging using excess of loss contracts should start with large losses and reduce the lower strike price as the amount spent on hedging increases (see, for example, Froot 2001). The initial hedge ratio is set equal to the company's market share in the state or region. (2) The evolutionary solver is then used to estimate the upper and lower strike prices and hedge ratios that minimize the criterion function. (3) Using the strike prices and hedge ratio from stage 2 as starting values, optimization is then conducted using the Newton-Raphson algorithm. (4) Stages 1 to 3 are then repeated for the next highest cost constraint using as starting values in stage 1 the stage 3 optimization results from the previous cost constraint. (5) After completing the optimization for all ten cost constraints, we perform a visual inspection of the resulting variance, VaR, or EEV reduction frontier to check for discontinuities and reestimate where necessary.

[23]Discontinuities were rare and represented cases where the optimization failed to find the global minimum.

To provide further information about the quality of the solutions, we conducted a randomization test using alternative starting values. The test was applied to twenty firms – five firms chosen randomly from each of the size quartiles. Using the variance reduction criterion function, the optimization problem was reestimated thirty times for each firm, and all cost constraints, for a total of 6,000 additional model solutions. In each of the reestimations, starting values for the striking prices and hedge ratios were randomly selected from intervals with lower bounds equal to 50 percent of the original starting values and upper bounds equal to 150 percent of the starting values that would have been used in our original five-step approach.[24] E.g., if the statewide hedge lower strike equals M, a starting value for a given trial was randomly selected from the interval [0.5M,1.5M], and the same procedure was followed for the upper strike and hedge ratio. Using the randomized starting values, we then apply the two-stage evolutionary genetic/Newton-Raphson approach to provide thirty sets of solutions for all ten cost constraints for each of the twenty firms.

The results of the analysis are robust to the randomization of starting values. The strike prices and hedge ratios for the thirty reestimations for each cost constraint tended to fluctuate in narrow bands around the original estimates. The solution value of the objective function was used as the primary statistic to evaluate the robustness of the results, i.e., we sought to determine whether randomization could substantially reduce the standard deviation of the insurer's hedged net loss in comparison with the original solution. More specifically, we computed the <u>maximum</u> percentage reduction in the original hedged net loss standard deviation achieved over the thirty runs for each cost constraint. The maximum reductions were then averaged across the ten cost constraints for each insurer. Based on these averages, the randomization results would have reduced the statewide net loss standard deviations by 0.12 percent and would have reduced the regional standard deviations by 0.53 percent.[25] We also evaluated the <u>maximum</u> reduction in standard

[24]That is, for each of the thirty randomizations for each firm, we followed the procedure outlined in footnote 22 except that starting values at each cost constraint were randomly selected in the [0.5,1.5] band around the starting values that would have been used in the non-randomized case.

[25]These averages are obtained by first averaging the maximum standard deviation reductions across the

deviation achieved for each company over its three hundred randomizations. The maximum reduction based on the statewide hedge was less than 1 percent for fifteen firms, less than 2 percent for nineteen firms, and less than 3 percent for all twenty firms. For the regional hedge, the maximum reduction was less than 1 percent for ten of the twenty firms, less than 3 percent for fifteen firms, less than 5 percent for eighteen firms, and between 5 and 8 percent for the remaining two firms.[26] Thus, the randomization results confirm the robustness of the original solutions. Additional fine tuning leads to only slight improvements in the objective function, and the randomized results support the same conclusions as the original solutions.[27]

3. The Basis Risk Analysis: Results

In this section, we present the results of our empirical analysis of hedging effectiveness using index-linked CAT securities. We begin the section by providing more details about the sample, the hurricane simulation results, and the loss indices, and then present the analysis of hedging effectiveness.

The Sample

As mentioned above, the study uses 1998 data on the value of residential property exposed to catastrophic loss in Florida, provided by the Florida Insurance Commissioner. The data base includes 255 of the 264 property insurers operating in Florida in that year, accounting for 93 percent of the insured residential property values in the state.[28] Thus, the study measures hedging effectiveness for residential

---

ten cost constraints for a given insurer and then averaging these results across insurers.

[26]As another way of viewing the results, the objective function value was reduced by 1 percent or more in only about 5 percent (309) of the 6,000 trials.

[27]As a further robustness check, we also conducted experiments where we reversed the order of the solution algorithms, first running Newton-Raphson and then running the genetic algorithm. As expected based on the general optimization literature (e.g., Pinter 1996), these results were much noisier and did not improve upon the results obtained using the procedure discussed in the text.

[28]The residential data include coverage under the following types of property insurance policies: apartment buildings, condominium associations, condominium unit owners, dwelling fire and allied lines, farmowners, homeowners, mobile homes, and tenants policies. Data were not available on commercial property exposures.

property insurance.[29] The total value of insured residential property in Florida in 1998 was $764 billion.

More details on the sample are provided in Table 1. The table shows that the distribution of exposures across the companies in the industry is highly skewed, with the top quartile of insurers accounting for 94.7 percent of insured exposure in our sample. This is important from a public policy perspective because larger insurers are expected to be able to hedge more efficiently than smaller firms. Thus, even though some individual firms may not be able to reduce risk significantly by trading in index-linked derivatives, a high proportion of the total exposure in the state is likely to be subject to effective hedging.

Larger firms tend to have their exposures dispersed across a wider range of counties than smaller firms, indicating better diversification. On average, firms in the top quartile have exposures in 58 of the 67 counties in Florida, compared with 44, 29, and 12 counties for insurers in the second, third, and fourth size quartiles. Larger firms also tend to be more diversified in terms of the coefficient variation of the market share across counties and in terms of the county market share Herfindahl index. These results suggest that larger firms should be able to hedge more efficiently using index-linked contracts than smaller insurers**.**

Simulation Results and CAT Loss Indices

The second step in the analysis is to simulate county-level losses for the insurers in the sample using the AIR model. We initially simulated 10,000 years of hurricane experience. In order to reduce the time required to perform the optimization analysis, we base most of the analysis on a random sample of 1,000 years of experience from the simulated 10,000 year data base. Robustness checks based on conducting the optimization using the full 10,000 years of experience for a random sample of 10 insurers revealed that virtually no accuracy is lost by basing most of the analysis on the 1,000 year random sample of events.

The simulations produce the variables L_{jkrt} = hurricane losses for company j, in county k_r, located in intra-state region r, for simulation year t, where j = 1, . . . , 255; k_1 = 1, . . . , K_1; . . . ; k_4 = 1, . . . , K_4; r =

[29]This is the type of insurance with the most significant catastrophic risk financing problem because business firms are better able to search the market for insurance coverage and have access to alternative hedging mechanisms such as captive insurance companies.

1, . . . , 4; t = 1, . . . , 10,000 (1,000 for most of the analysis), k_r = the kth county in region r, and K_r = the number of counties in region r. The simulated losses are then used to construct the following loss indices:

$$\text{The "Perfect" Index} = L^P_{j \cdot \cdot t} = \sum_{r=1}^{R} \sum_{k_r=1}^{K_r} L_{jk_rrt} \tag{10}$$

$$\text{The Regional Indices} = L^R_{\cdot \cdot rt} = \sum_{j=1}^{N} \sum_{k_r=1}^{K_r} L_{jk_rrt} \tag{11}$$

$$\text{The State Index} = L^S_t = L_{\cdots t} = \sum_{j=1}^{N} \sum_{r=1}^{R} \sum_{k_r=1}^{K_r} L_{jk_rrt} \tag{12}$$

where N = the number of insurers (255), R = the number of regions (4), K_r = the number of counties in region r, and a dot in place of a subscript means that a summation has been taken over that subscript. Hedge portfolios are formed for each insurer to determine the basis risk for each index.

Non-Linear (Call Spread) Hedging

As mentioned above, the analysis focuses on non-linear, "large loss" hedging strategies. Insurers are assumed to form hedge portfolios consisting of their own losses and a position in call option spreads on loss indices. The hedge ratios and option strike prices are then chosen to minimize a criterion function subject to a cost constraint. The risk measures minimized are the variance, the value at risk (VaR), and the expected exceedence value (EEV) of the insurer's net loss liabilities, where net loss liabilities are defined as unhedged loss liabilities minus the payoff on the hedge. Using the same hedging criteria, we also analyze hedging strategies that focus on the index-hedge basis, defined in equation (4) as the difference between the payoff on the index hedge and the payoff on the perfect hedge.

The cost constraints are specified as percentages of the insurer's expected Florida homeowners losses, ranging from 5 percent to 50 percent. By varying the cost constraint, an efficient frontier is generated based on each of the criteria. To focus purely on basis risk, most of the analysis is conducted under the assumption that hedging contracts are available at prices equal to the expected losses under the contracts, i.e., the expected recovery from the hedge. We also report robustness tests based on the assumption that the options are priced at the expected loss plus a risk premium.

We first consider the effect of hedging on the variance of the insurer's net loss. Figure 2 shows the variance-reduction frontiers for the insurers in the largest size quartile, obtained by varying the cost constraint, on the assumption that insurers hedge to protect against losses from a $1 billion statewide loss event.[30] Each point on the frontier is an unweighted average of the percentage variance reduction across the firms in the top quartile for the specified cost constraint. The figure compares frontiers based on the perfect hedge, the state hedge, and the regional hedge. The results confirm that hedging with the regional loss indices is more effective than hedging using the state loss index. In fact, the variance reduction using the regional hedge is closer to that given by the perfect hedge than to the variance reduction based on the statewide hedge. For example, an expenditure of 15 percent of expected losses reduces the conditional net loss variance by 40 percent using the statewide hedge, 58 percent using the regional hedge, and 62 percent using the perfect hedge. Thus, the basis risk of the regional hedge is not very large and might be worth incurring in order to avoid the moral hazard inherent in the perfect hedge.

The average variance reduction frontiers for insurers in the four size quartiles based on the regional hedge are shown in Figure 3, again based on the $1 billion statewide loss threshold. Perhaps the most surprising result is that the frontiers in the two largest size quartiles are almost indistinguishable. Thus, the insurers in the top two quartiles can hedge with about equal effectiveness using the regional loss indices, and the quartile 3 results are almost as good. This suggests that it is not size per se but rather diversification that determines hedging effectiveness, at least for insurers in the top three size quartiles. As expected, the degree of variance reduction is noticeably less for insurers in the fourth size quartile.

To provide additional information on basis risk for the sample insurers, Figure 4 shows the frequency distribution of the conditional variance-reduction hedge efficiency in the $1 billion layer for an expenditure

---

[30]The results based on the $2.5 and $5 billion thresholds are similar and therefore not shown.

of 15 percent of expected losses.[31] The most striking result based on Figure 4 is that the regional hedge is at least 95 percent as effective as the perfect hedge in terms of reducing conditional loss volatility for 76 of the 255 firms in the sample and at least 90 percent as effective as the perfect hedge for 143 firms. These results provide further evidence that the degree of basis risk from index hedging may be sufficiently small to make index hedging attractive for most Florida insurers. The statewide hedge is at least 90 percent as effective as the perfect hedge for 36 of the 255 firms, again at the 15 percent expenditure level with the $1 billion loss threshold. However, the statewide hedge is no more than 50 percent efficient for 105 of the sample firms, confirming that many firms cannot hedge efficiently based on the statewide index.[32]

From a public policy perspective, it is relevant to consider the proportion of total exposures in the state that could be hedged efficiently using index-linked contracts. These results are shown in Figure 5, for an expenditure of 15 percent of expected losses based on the $1 billion industry loss threshold. Reflecting the skewness in exposures across firms in the industry and the relatively high diversification of the largest firms, Figure 5 reveals that 70 percent of the exposures in our sample could be hedged with at least 95 percent efficiency and 92 percent could be hedged with at least 90 percent efficiency using regional index contracts. Even using the less-effective state-wide loss index, about 36 percent of the total exposures could be hedged with at least 95 percent efficiency and 55 percent could be hedged with at least 90 percent efficiency. Thus, the development of a robust market for index-linked Florida CAT call spreads, especially based on regional indices, would provide an effective solution to the state's CAT loss financing problem.

[31]The results for other expenditure levels are comparable and thus not shown. Recall that hedge efficiency is defined for the variance reduction criterion as the ratio of the variance reduction using the statewide and regional hedges to the variance reduction under the perfect hedge (see equations (8) and (9)).

[32]We also have developed a parametric index, based on a regression model with the natural log of the by regressing the dollar value of simulated losses from storms on three physical measures of storm severity – the natural logs of 30 minus the central pressure at the eye of the storm, the forward velocity of the storm, and the radius to maximum wind speed. The analysis shows that insurers can hedge almost as effectively using the parametric index (the predicted values from the regression model) as they can using the actual state and regional loss indices. The results of the parametric analysis are available from the authors.

We next consider the value at risk (VaR) and expected exceedence value (EEV) hedging criteria. Because the analyses of these two criteria lead to the similar conclusions and the EEV has more desirable theoretical properties than the VaR (Artzner, et al. 1999), we focus the discussion on the EEV.[33] Recall that the EEV is the expected loss, conditional on losses exceeding a specified threshold (see equation (7)). We report on the EEV based on the 95th percentile of each insurer's unhedged loss distribution, i.e., the EEV minimized above the $VaR_j(0.05,L_j)$. Tests at the 92.5th and 97.5th percentiles yielded similar conclusions. The ratio of the expected CAT loss above the 95th percentile to the total expected CAT loss ranges monotonically from 19 percent for firms in the first quartile to 31 percent for firms in the fourth quartile. Thus, hedges based on this threshold also have the potential to significantly reduce the insurers' overall expected CAT losses.

The expected exceedence value (EEV) reduction frontiers for the firms in the largest size quartile are shown in Figure 6.[34] The results again support the conclusion that insurers in the top size quartile can hedge effectively using the regional loss indices. For example, a 50 percent reduction in the EEV can be obtained at a cost of about 15 percent of expected losses with the perfect hedge and about 16 percent of expected losses for the regional index hedge. A comparable reduction costs about 21 percent of expected losses under the statewide hedge.

The regional and perfect hedge frontiers diverge significantly for the higher expenditure levels in comparison with the variance reduction frontiers. For example, at an expenditure of 30 percent of expected losses, the EEV is reduced by 80 percent using the regional hedge and by 92 percent using the perfect hedge. Hence, the efficiency of the EEV hedges tends to be lower than the efficiency of variance hedges at the

[33]The VaR results are available from the authors.

[34]The EEV using the perfect hedge is reduced to zero at an expenditure of approximately 40 percent of the expected loss for all insurers in the top size quartile because the unhedged EEV for most top quartile insurers is less than 40 percent of expected loss. Recall that hedging at actuarially fair prices will not reduce the expected loss for the firm -- it will only change the distribution of those losses.

higher expenditure levels. Nevertheless, at the 20 percent expenditure level, 97 of the firms in the sample can hedge with at least 95 percent EEV efficiency using regional contracts and 138 can hedge with at least 90 percent efficiency. Also generally comparable to the variance results, 66 percent of exposures in the sample can be hedged with at least 95 percent EEV efficiency and 81 percent can be hedged with at least 90 percent EEV efficiency, again at the 20 percent expenditure level. Thus, the EEV results generally tend to confirm the results based on variance reduction.

Finally, we measure hedge effectiveness when the variable optimized is the index-hedge basis, defined as the difference between the amount collected on the index hedge and the amount collected on the perfect hedge (see equation (4)). The primary reason for considering the basis is to address the concerns of insurers that the index hedge will pay off significantly less than the perfect hedge in the event of a large loss. The analysis of the basis focuses on the minimizing its variance, conditional on receiving a payment on the perfect hedge. That is, the objective function is: $\sigma_T^2[B_j^i(h_j^i,M_j^i,U_j^i) * L_j \, 0 \, L^T , L_j^P \ldots L_j]$ = the variance of the basis for insurer j using index i (regional or state), and the payoff L_j^P is based upon the variance minimizing perfect hedge conditional on the \$1 billion, \$2.5 billion, or \$5.0 billion industry-wide loss threshold. The condition $L_j^P \ldots L_j$ is added here because the objective is to minimize the risk that the payoff on the index hedge departs from the payoff on the perfect hedge, given that there is a payoff on the latter hedge. Variance minimization is emphasized because we can then provide an intuitive probabilistic interpretation of the likely difference between the index payment and the perfect payment using the standard deviation of the basis.

In the index-hedge basis analysis, we minimize the variance of the basis subject to the usual cost constraints. However, in this minimization, we found that the minimum variance solution often resulted in an estimated *average* basis significantly different from zero. Accordingly, we also impose the constraint that the average basis be within plus or minus 1 percent around zero and, alternatively, plus or minus 5 percent from zero, where percentages were obtained by dividing the basis by the company's gross loss from an event.

The results of the minimum variance basis analysis show that the variance reduction frontier and efficiency of the basis optimized results are very similar to those based on the conditional net loss variance optimizations. Because it is the basis and not the insurer's net loss that is the subject of the optimization, the variance reduction and hedging efficiency are slightly less in the conditional basis minimization results than in the conditional variance minimization results.

The more important findings for the basis variance minimization results, however, concern the magnitude of the standard deviation of the basis. These results are presented in Figure 7, which shows the average standard deviation of the basis for the firms in the largest size quartile.[35] The figure shows the standard deviation for hedges based on minimizing the conditional basis variance and, for comparison purposes, the standard deviation of the basis resulting from the conditional variance minimizing hedges. The results are for the $1 billion industry-wide loss threshold and, in the basis optimized case, with the ± 5 percent average basis constraint. In each case, regional contracts are used to construct the hedge.

Figure 7 shows that the standard deviation of the basis is about 10 percent for the 5 percent cost constraint and then rises to about 25 percent for the 50 percent cost constraint. The basis standard deviation is higher for the variance minimized hedges because these hedges were not constructed by minimizing the basis variance but rather the conditional variance of the net loss. Also shown on the chart are the 10 percent values at risk for the basis optimized and variance optimized cases. These curves show the percentage shortfalls of the index hedge collection in comparison with the perfect hedge collection that are likely to occur with a probability of 10 percent. These values were not obtained by minimizing the VaR of the basis but rather are the implied VaRs based on hedges obtained from the basis variance optimization and conditional variance optimization exercises.

There are two main points to be made based on Figure 7. First, the basis standard deviations and the absolute values of the 10% VaRs are positively related to the cost constraint. The reason for this relationship

[35]The results for firms in the second and third size quartiles are similar and thus not shown. The standard deviations and VaR index-hedge basis statistics are larger (in absolute value) for firms in the fourth size quartile.

is that the lower cost constraints devote resources to hedging the largest losses, which tend to produce the highest correlations between the industry losses and the individual insurer losses, leading to lower basis risk. Second, the standard deviations and VaR are fairly large, especially for the larger cost constraints. For example, with a 10 percent cost constraint, the basis optimized standard deviation is about 13 percent and the 10 percent VaR is -20 percent. This implies that there is a 10 percent chance of collecting 80 percent or less of what would have been collected under the perfect hedge for the relatively large losses that are hedged with the 10 percent cost constraint. For the 20 percent cost constraint, the 10 percent VaR is about 25 percent and the VaR increases (in absolute value) for higher cost constraints.

Hedge Effectiveness and Conditional Hedging: Further Discussion

Hedge Effectiveness. The results suggest that the effectiveness of index-linked contracts in hedging catastrophic risk depends upon the objectives of the firm and its risk-tolerance. If the objective is to reduce the conditional variance of the insurer's net losses (or the conditional VaR or EEV), then our results show that many insurers can hedge very efficiently using regional hedge contracts. Hedging with ninety or 95 percent efficiency relative to the perfect hedge would seem to satisfy the objectives of most decision makers, especially if they were encountering coverage availability problems or high prices in the reinsurance market.

The conclusion is less straightforward if the objective of the hedger is to avoid outcomes involving significant under-collections on large events in comparison with the perfect hedge. For example, from Figure 7, there is a 10 percent chance of under-collecting by 15 to 25 percent if the cost constraint is in the range that leads to hedging of the most severe losses (cost constraints of 5-20 percent). The attractiveness or lack of attractiveness of such hedges ultimately must be determined in a market context. However, there are several reasons to believe that such hedges may be economically beneficial relative to reinsurance.

First, even though we have been referring to reinsurance and hedger-specific CAT bonds as "perfect" hedges, in reality neither type of contract actually is perfect. Both contracts are usually sold with co-payment provisions, primarily to control moral hazard. For example, catastrophe excess of loss reinsurance tends to be sold with co-payments ranging from 5 to 20 percent, and the USAA CAT bond issues have contained a co-payment of 10 or 20 percent, depending on the issue (Froot 2001). Because the USAA CAT bond issues have been fully collateralized, these provisions are equivalent to having a 100 percent probability of collecting 80 or 90 percent of the loss in the call spread coverage layer. This compares to a 90 percent probability of collecting at least 75 to 85 percent with the index options portrayed in Figure 7.

Second, index-linked contracts have the potential to trade at significantly lower margins above the expected loss than insurer-specific contracts because they are less affected by moral hazard and have the potential to be more liquid. If prices are sufficiently low for index-linked contracts, it may make these contracts attractive to hedgers even in the presence of some basis risk.

Third, the value of the index contracts is primarily in softening the blow of a large catastrophic event. Recovering 75 to 90 percent of a large catastrophic loss in the hedged layer may place the uncovered loss in a range of magnitude that may be sustainable without substantial disruption to the firm's operations, i.e., even a partial recovery may reduce the net loss to a size category where expensive purchased hedges are less desirable. Fourth, in the case of reinsurance (although not for fully collateralized CAT bonds), the hedger has to be concerned about the credit risk of the reinsurer. Thus, the probability of full payment under reinsurance contracts is also less than 100 percent.[36] And, finally, although insurers have traditionally been accustomed to thinking in terms of hedges that reduce risk to zero, in more general financial theory and practice it is usually neither necessary nor desirable for hedging to totally eliminate risk. Thus, index linked catastrophe hedges are likely to appear more attractive to insurers as they begin to view hedging from the perspective of financial management rather than traditional insurance management.

Conditional and Unconditional Hedging. The preceding analysis has emphasized conditional hedging. That is, we formed hedges on the assumption that insurers hedge large losses, defined as losses

[36]Insurers also are likely to face counterparty risk in the market for index-linked CAT securities. However, properly designed margin and collateral requirements can be used to significantly reduce counterparty risk for exchange traded contracts.

arising from an industry-wide event that causes $1, $2.5, or $5 billion in property damage state-wide.

As a robustness check, we also conducted the analysis on an *unconditional* basis. That is, rather than forming hedges by solving optimization problems where insurers were constrained to minimizing risk above some conditioning point, we solved the hedging problem without imposing the conditioning constraint, potentially allowing for hedges to be established to reduce risk attributable to smaller losses. We then computed the relative efficiency of the optimal hedges based on all four conditioning thresholds (zero, $1 billion, $2.5 billion, and $5 billion) in minimizing variance of net losses above the $1 billion industry loss.

The comparison showed that the conditional variance above $1 billion is essentially the same for each of the four conditioning thresholds. The reason for this outcome is that the optimization process expends resources such that each additional dollar spent on hedging has the maximal marginal benefit in terms of reducing the criterion functions. Thus, even with no conditioning threshold, the optimization led to the hedging of the largest losses first, followed by the allocation of additional hedging expenditures to progressively smaller losses. We did not observe solutions where funds were expended to hedge relatively small losses while larger losses remained unhedged. Thus, our optimization process and results are consistent with theories of optimal reinsurance (e.g., Froot 2001), which also suggest that the largest events have the highest hedging priority.[37]

Hedging at Market Prices

The analysis so far has been conducted under the assumption that call spread contracts are available at actuarially fair prices equal to the expected loss under the contracts. The rationale for this approach is that

[37]As pointed out in Froot (2001) and Swiss Re (1997), actual reinsurance purchases generally do not conform to the theory, i.e., some insurers appear to hedge relatively small CAT losses and leave larger ones under-reinsured. However, this departure is usually explained by the presence of market imperfections which lead insurers to adopt strategies that are less than optimal due to price markups, availability problems, and other friction costs. In this part of the paper, we are investigating the basis risk insurers would encounter in buying actuarially fair CAT loss securities in markets without significant imperfections or friction costs. Hence, it is not surprising that our results are consistent with theory rather than with observed practice. We discuss non-actuarial pricing below.

catastrophic loss contracts are expected to be priced close to their actuarial value in informationally efficient, liquid securities markets, provided that catastrophic losses have low systematic risk.[38] However, because most catastrophic risk derivatives issued to date have been sold at prices in excess of the expected actuarial losses, we also conduct our hedging analysis under the assumption that CAT security prices are actuarially unfair. We base the analysis on the recent prices for CAT bonds and all available prices for CBOT Florida call spreads shown in Appendix Table A.1.

The contractual forms in the non-actuarial analysis are identical to those used in the hedging analysis above, the only difference being that the contracts analyzed in this section are priced at a markup over the expected loss. The perfect hedge contracts are analogous to insurer-specific CAT bonds or reinsurance, whereas the index hedge contracts are analogous to CBOT options. Accordingly, the perfect hedge contracts are assumed to be sold at a premium-to-expected-loss ratio of 6.8 and the index hedge contracts are assumed to be sold at a premium-to-loss ratio of 2.1, matching the median risk premia in Appendix Table A.1.

The results of the non-actuarial hedging analysis are shown in Table 2. Because the results under different hedging strategies lead to the same conclusions, only the variance reduction results are shown. Table 2 shows the ratios of hedge effectiveness using market price contracts to the hedge effectiveness that could be achieved using actuarially priced perfect hedge contracts, for each of the ten cost constraints used in our analysis.[39] The ratios in the table are averages based on a stratified (by size quartile) sample of the firms in our data base. The sample consists of three firms chosen randomly from each size quartile.

The results in Table 2 show that insurers can still significantly reduce their conditional variances using index hedging, even when option pricing is non-actuarial. However, as expected, hedge effectiveness

---

[38]Evidence that catastrophic risk contracts have low systematic risk is presented in Litzenberger, Beaglehole, and Reynolds (1996).

[39]We do not show the results based on a common mark-up for all securities because the relationships between the perfect hedge contracts and the index hedge contracts would be the same as in the actuarially fair cases discussed above. I.e., the efficiencies shift downward, but the relative performance remains the same.

is reduced in comparison with actuarially fair perfect hedge contracts. For example, if expenditures on hedging are constrained to 25 percent of expected losses, the market priced perfect hedge reduces the conditional variance by only 32.8 percent of the perfect hedge variance reduction that could be obtained with actuarial prices. The results with the state and regional hedges are better because the markup over the actuarial price is significantly less than for the perfect hedge contracts. With the 25 percent cost constraint, the market price state hedge reduces the conditional variance by 51.8 percent of the reduction that could be achieved using actuarially priced contracts, and the comparable reduction for the regional hedge is 63.7 percent. Thus, market pricing reduces hedge efficiency and markups can significantly change the relative efficiency rankings among the alternative indices.

The size of the markup over expected losses is obviously critical in determining the hedging efficiency of insurance derivative contracts. Such contracts must compete with excess of loss reinsurance - the traditional hedge for insurers facing CAT loss exposure. Interestingly, the markups on the insurance derivative contracts shown in Table A.1 are broadly consistent with markups on catastrophe reinsurance contracts. Froot and O'Connell (1999) show that price-to-loss ratios during the late 1980s and early 1990s for excess of loss property reinsurance contracts ranged from about 1.5 in 1987, to 3.0 in 1992, and to 7.0 in 1994, all in the same range as the price-to-loss ratios in Table A.1. Thus, CAT securities may be price-competitive with reinsurance even with the relatively high markups in today's CAT securities market.

The price-to-expected-loss ratios on insurance derivatives can be expected to decline relative to reinsurance as the market becomes more mature. Reinsurance is sold by firms that have limited capital and are averse to insolvency risk; whereas CAT loss derivatives are closer to being pure financial instruments, not dependent upon the solvency or capitalization of any specific firm or industry.[40] Consequently, CAT loss securities are more likely to approach actuarial fairness than reinsurance, particularly for mega-CATs that

[40]The imperfections of the reinsurance market are further discussed in Berger, Cummins, and Tennyson 1992, Froot and O'Connell 1999, Cummins and Weiss (2000), and Froot 2001.

would significantly stress the capacity of world insurance markets. A significant reduction in spreads for securitized CAT instruments would likely lead to the dominance of reinsurance by securitized instruments for CAT risk finance even in the presence of the moderate basis risk.

4. Conclusions

The securities market has responded to the dramatic increase in catastrophic (CAT) property losses over the last decade by developing innovative new derivative securities to finance this type of loss. The introduction of catastrophe loss securities also has been driven by the increasing recognition that the insurance and reinsurance markets do not provide efficient mechanisms for financing losses from low frequency, high severity events.

CAT loss securities have been structured to pay-off on three types of variables – insurer-specific catastrophe losses, insurance-industry catastrophe loss indices, and parametric indices based on the physical characteristics of catastrophic events. Securities based on insurer-specific losses have no basis risk but expose investors to moral hazard; whereas securities based on industry loss indices or parametric triggers greatly reduce or eliminate moral hazard but expose hedgers to basis risk. Loss-index contracts are easier to standardize than insurer-specific contracts, potentially giving them lower transactions costs and higher liquidity. Thus, such contracts would be likely to dominate insurer-specific contracts, provided that basis risk is sufficiently low. This paper provides a comprehensive analysis of the basis risk insurers would encounter in hedging catastrophic risk in Florida, the state with the highest exposure to windstorm risk.

The analysis begins by simulating 10,000 years of hurricane losses in Florida and measuring the effects of the resulting catastrophes on 255 insurers representing 93 percent of the residential property exposed to loss in the state. We then measure basis risk and hedge efficiency for the sample insurers using a variety of hedging strategies.

In our hedging analysis, we form portfolios consisting of a short position in insurer loss liabilities and a long position in call option spreads on loss indices. Three indices are analyzed – a “perfect” index consisting of the insurer's own losses, a statewide industry loss index, and four intra-state regional industry loss indices. Three criterion functions are minimized, subject to cost constraints – the variance of the insurer's hedged net losses, the value at risk (VaR), and the expected exceedence value (EEV), defined as the expected catastrophic loss conditional on the loss exceeding a specified threshold. We measure *hedge efficiency* as the ratio of the risk reduction obtained using loss index options to the risk reduction obtained using the perfect index. Efficiency is also evaluated using the standard deviation and the 10 percent VaR of the *index-hedge basis*, defined as the amount collected on an index hedge minus the amount collected on the perfect hedge.

The hedging problems are solved under the assumption that insurers seek to hedge relatively large losses, consistent with observed insurer behavior in the CAT bond market. Most of the minimization problems are solved conditional on losses arising from events causing industry-wide losses of at least $1, $2.5, or $5 billion.

The principal finding is that firms in the three largest Florida market-share quartiles can hedge almost as effectively using intra-state index contracts as they can using perfect-hedge contracts, when the objective is to minimize the conditional variance, VaR, or EEV. For example, the hedges based on regional index contracts are at least 90 percent as effective as the perfect hedge in terms of reducing loss volatility for 143 of the 255 firms in the sample and at least 95 percent as effective for 76 of the 255 sample firms. Hedging with the statewide contracts, on the other hand, is efficient only for insurers with the largest state market shares and insurers that are highly diversified throughout the state. Thus, the regional contracts hold significant promise for the development of a more liquid market for index-linked CAT securities. Hedging with regional contracts also offers insurers and policy makers a solution to the catastrophic risk financing problem in Florida because 93 percent of the property exposure in the state could be hedged with at least 90 percent efficiency using the regional contracts.

Analysis of the index-hedge basis reveals that insurers would face a 10 percent probability of collecting 12 percent less than the loss amount for the lowest cost constraint, ranging upward to 30 percent for the highest cost constraint. In considering the implications of this result for index hedge effectiveness, it should be kept in mind that the basis VaR is lowest for hedges involving the largest losses, that CAT reinsurance usually involves a co-payment of 5 to 20 percent, and that buying reinsurance exposes the hedging insurer to the reinsurer's credit risk. Thus, hedging through index contracts is likely to be attractive relative to reinsurance even with the degree of basis risk found in the index-hedge basis analysis.

As expected, hedging with contracts that are sold at mark-ups over the expected loss is less efficient than hedging using contracts sold at actuarially fair prices. Even at the current markups in the CAT securities market, however, insurance-linked securities generally are competitive with reinsurance in terms of price and hedge efficiency. Moreover, mark-ups in the CAT securities market can be expected to decline as investors acquire more experience with these contracts and the market becomes more liquid. CAT securities could come to dominate reinsurance for hedging CAT losses if prices converge towards actuarial fairness.

Overall, our analysis suggests that insurance-linked securities based on exchange-traded, index-linked contracts could be used effectively by insurers in hedging catastrophic risk. This is important given the inefficiency of the reinsurance market in dealing with this type of loss. To the extent that basis risk is perceived as too high by some potential hedgers, intermediaries such as reinsurers could solve the problem by forming diversified portfolios of primary insurers and hedging the residual risk in the CAT securities market. Hedging of catastrophic risk has the potential to avoid the destabilization of insurance markets resulting from a major event; and with more widespread trading, insurance-linked securities would play a price-discovery role, potentially smoothing the reinsurance underwriting cycle. The more widespread trading of insurance-linked securities would allow investors to shift the efficient frontier in a favorable direction by further diversifying their portfolios using these low-beta assets. Thus, the securitization of CAT risk has the potential to increase efficiency in both insurance and securities markets.

Table 1
Summary Statistics 1998 Florida Insurer Exposure Database

Variable	Size Quartile	Average	Std. Deviation	Minimum	Maximum
Statewide Exposure Limits	1	10,488,076,940	27,023,882,691	947,613,000	197,123,513,015
	2	489,399,825	209,101,097	212,101,944	917,368,990
	3	87,212,264	55,098,625	21,396,000	206,663,000
	4	6,603,762	7,059,451	1,000	21,090,000
	All Insurers	2,778,651,509	14,183,583,447	1,000	197,123,513,015
Statewide Market Share	1	1.373%	3.538%	0.124%	25.810%
	2	0.064%	0.027%	0.028%	0.120%
	3	0.011%	0.007%	0.003%	0.027%
	4	0.001%	0.001%	0.000%	0.003%
	All Insurers	0.364%	1.857%	0.000%	25.810%
Number of Counties with Exposure	1	58.344	11.360	15.000	67.000
	2	44.234	14.777	9.000	67.000
	3	29.203	19.145	3.000	67.000
	4	12.476	16.264	1.000	67.000
	All Insurers	36.157	23.095	1.000	67.000
% of Counties with Ocean Exposure	1	47.104%	9.248%	25.000%	100.000%
	2	52.657%	8.741%	38.636%	81.818%
	3	60.400%	17.039%	26.471%	100.000%
	4	70.612%	26.259%	0.000%	100.000%
	All Insurers	57.642%	18.931%	0.000%	100.000%
% of Exposures in Ocean Counties	1	70.150%	16.715%	23.198%	100.000%
	2	71.446%	18.197%	18.563%	99.657%
	3	70.092%	27.229%	8.702%	100.000%
	4	73.570%	31.800%	0.000%	100.000%
	All Insurers	71.306%	24.169%	0.000%	100.000%
County Market Share CoV	1	1.365	0.607	0.363	3.414
	2	2.204	1.143	0.720	5.983
	3	3.353	1.515	0.931	7.765
	4	5.380	2.165	1.316	8.185
	All Insurers	3.066	2.096	0.363	8.185
County Market Share Herfindahl	1	0.084	0.055	0.024	0.262
	2	0.126	0.116	0.030	0.649
	3	0.240	0.203	0.025	0.892
	4	0.448	0.315	0.035	1.000
	All Insurers	0.224	0.242	0.024	1.000

Note - Data obtained from Florida Department of Insurance regulatory filings. 264 insurer have exposure to losses due to hurricanes of which 255 insurers have usable data. The data set includes 92.8 percent of exposures in Florida subject to windstorm loss. Size quartiles based upon total residential property exposed to loss in Florida. Insurers in quartile 1 had 87.9% of exposure limits in the state. Quartiles 2, 3, and 4 had 4.1%, 0.73% and 0.054% of the exposure limits in the state, respectively.

Table 2
Conditional Variance Efficiencies:
Market Price Hedge Relative to Fair Price Perfect Hedge

Cost	Market/Actuarial		
% of E(Loss)	Perfect	Statewide	Regional
5%	0.1036	0.4416	0.5619
10%	0.2016	0.4740	0.5840
15%	0.2710	0.4863	0.5991
20%	0.3064	0.4983	0.6193
25%	0.3276	0.5182	0.6365
30%	0.3477	0.5321	0.6535
35%	0.3672	0.5511	0.6726
40%	0.3866	0.5666	0.6926
45%	0.4062	0.5827	0.7097
50%	0.4259	0.5953	0.7266

Note: The efficiencies are the ratios of conditional variance reduction above $1billion industry-wide loss events using hedge contracts with median market risk premia divided by the reduction obtained using the perfect hedge priced at the expected loss. The price-to-expected loss ratio for the perfect hedge contracts is 6.8, and the price-to-expected-loss ratio for the state and regional index contracts is 2.1. These ratios are the median ratios for the CAT bond and CBOT option contracts, respectively, shown in Appendix Table A.1.

Figure 1
Simulating Insured Losses Using the AIR Model

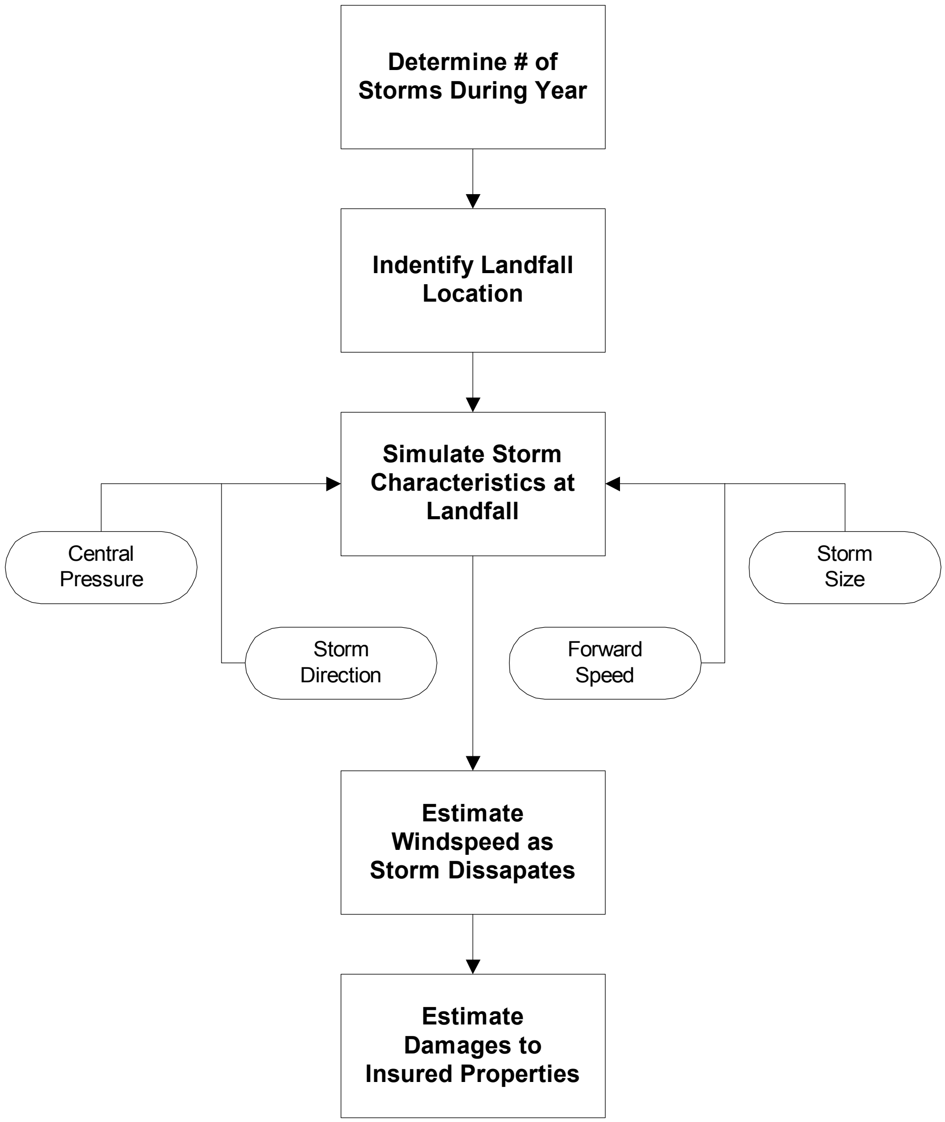

Figure 2: Conditional Variance Reduction Frontiers
Average for Insurers in Largest Size Quartile
$1Billion Industry-wide Loss Threshold

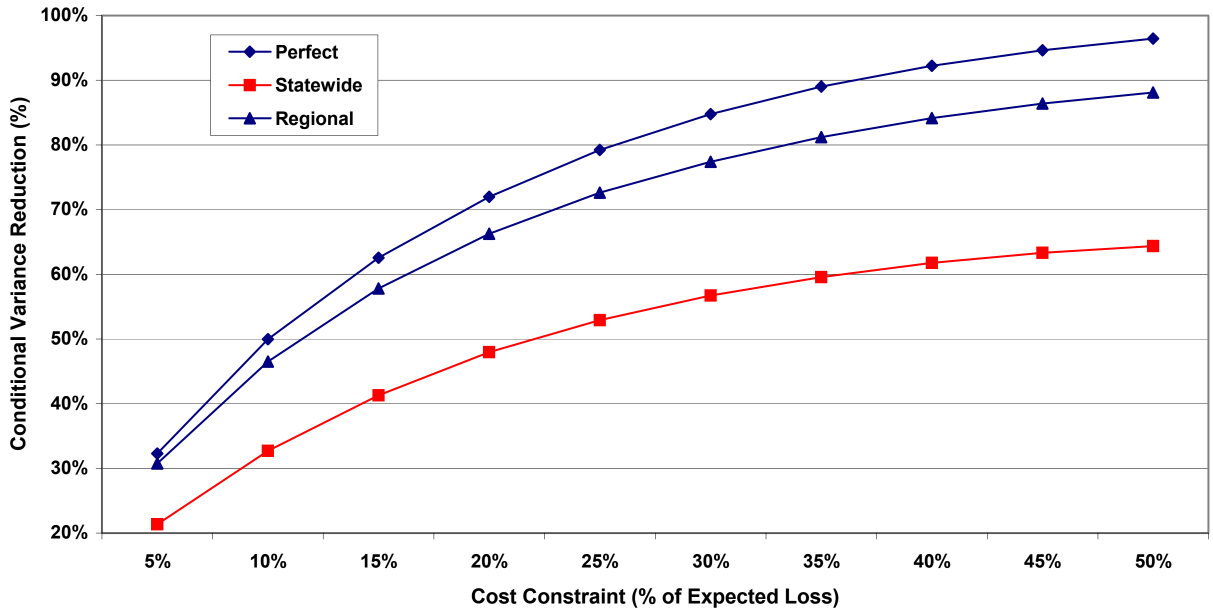

Figure 3: Conditional Variance Reduction Frontiers
By Insurer Size Quartile Using Regional Indicies;
$1Billion Industy-wide Loss Threshold

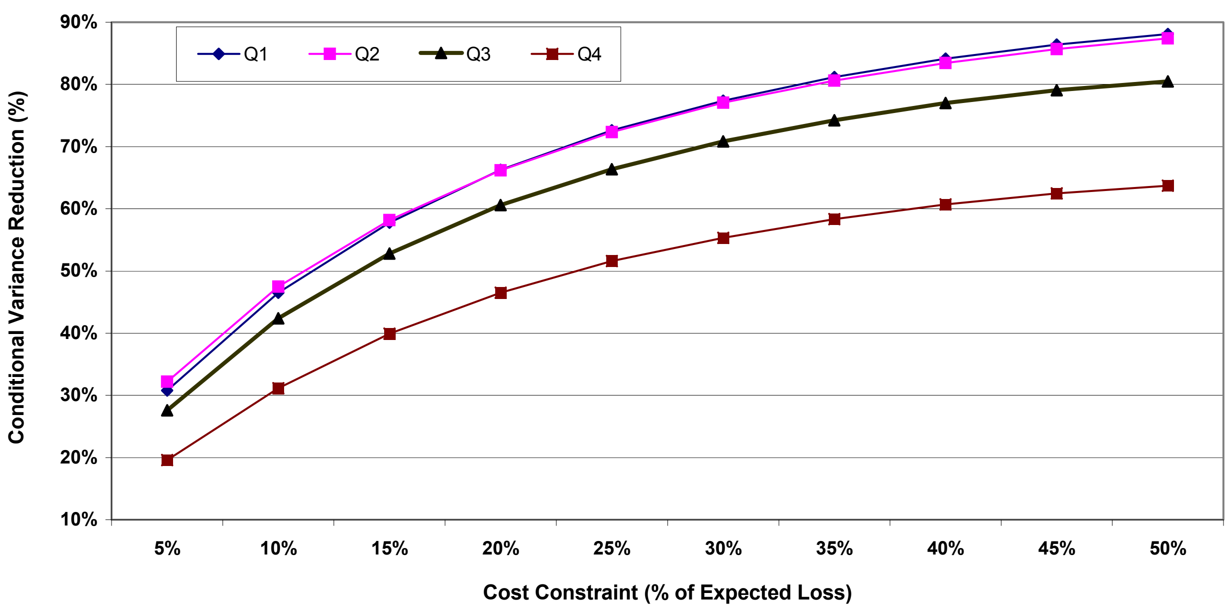

**Figure 4: Conditional Variance Reduction Hedge Efficiency By No. of Firms
Hedging Cost Constaint = 15% of Expected Loss;
$1Billion Industry-wide Loss Threshold**

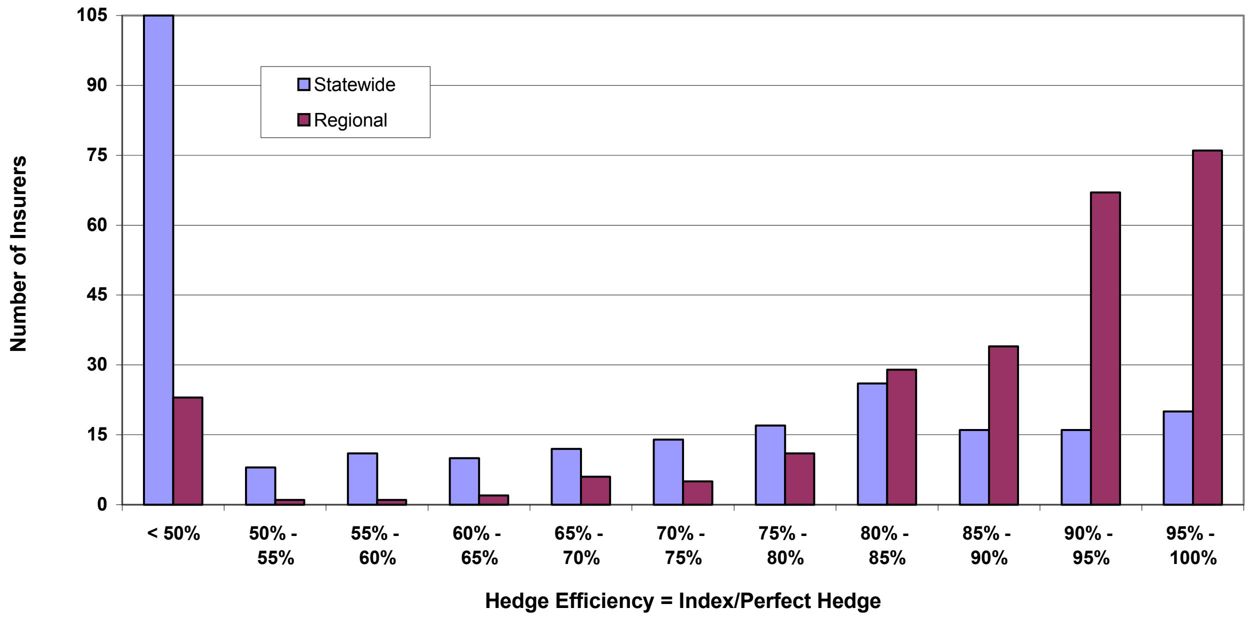

**Figure 5: Conditional Variance Reduction Hedge Efficiency By Exposures
Hedging Cost Constaint = 15% of Expected Loss
$1Billion Industry-wide Loss Threshold**

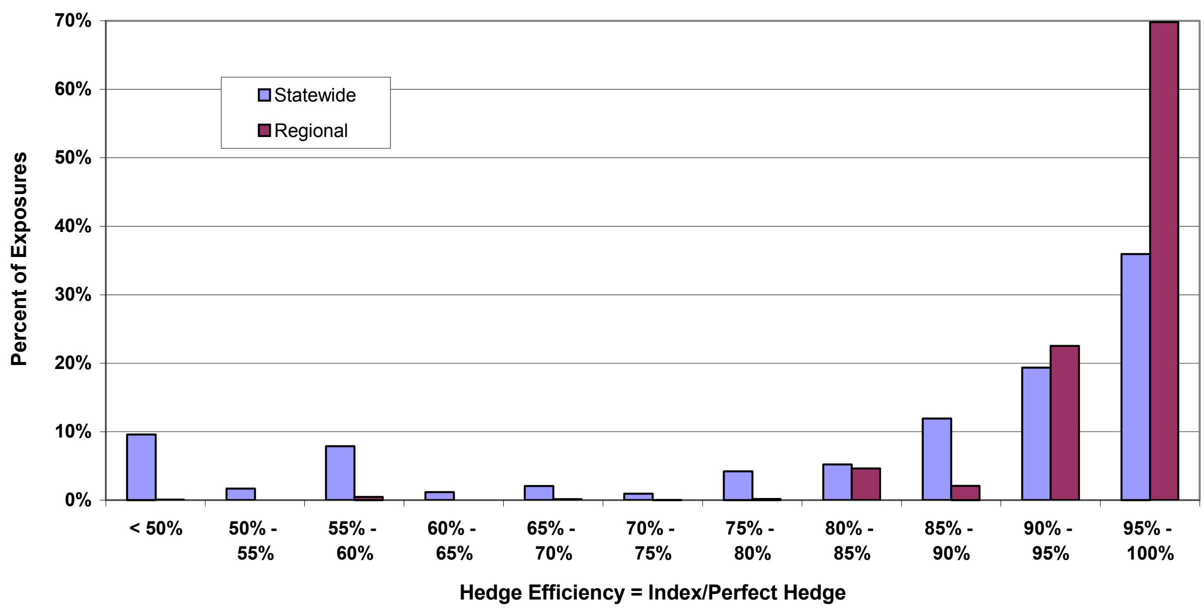

Figure 6: Expected Exceedance Value Reduction Frontiers
Average for Insurers in Largest Size Quartile
EEV Threshold = 95th Percentile

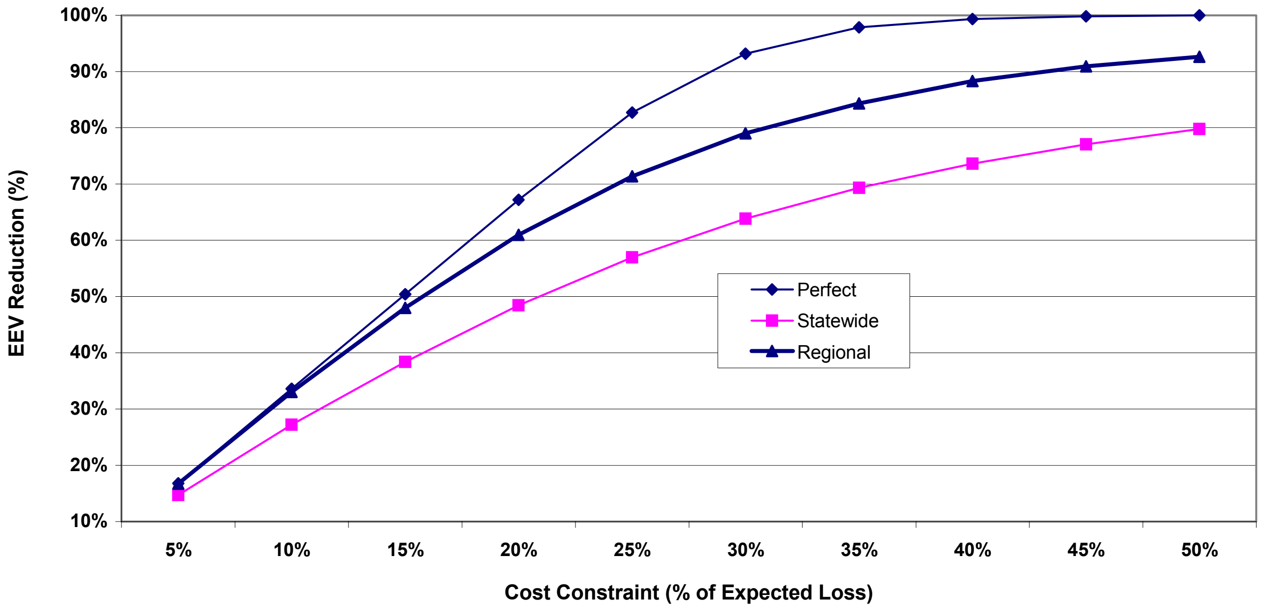

Figure 7: Standard Deviation and 10% VaR of the Basis :
Variance Optimized Vs. Basis Optimized Hedges
Average for Insurers in Largest Size Quartile
$1 Billion Threshold Industry-wide Loss Threshold

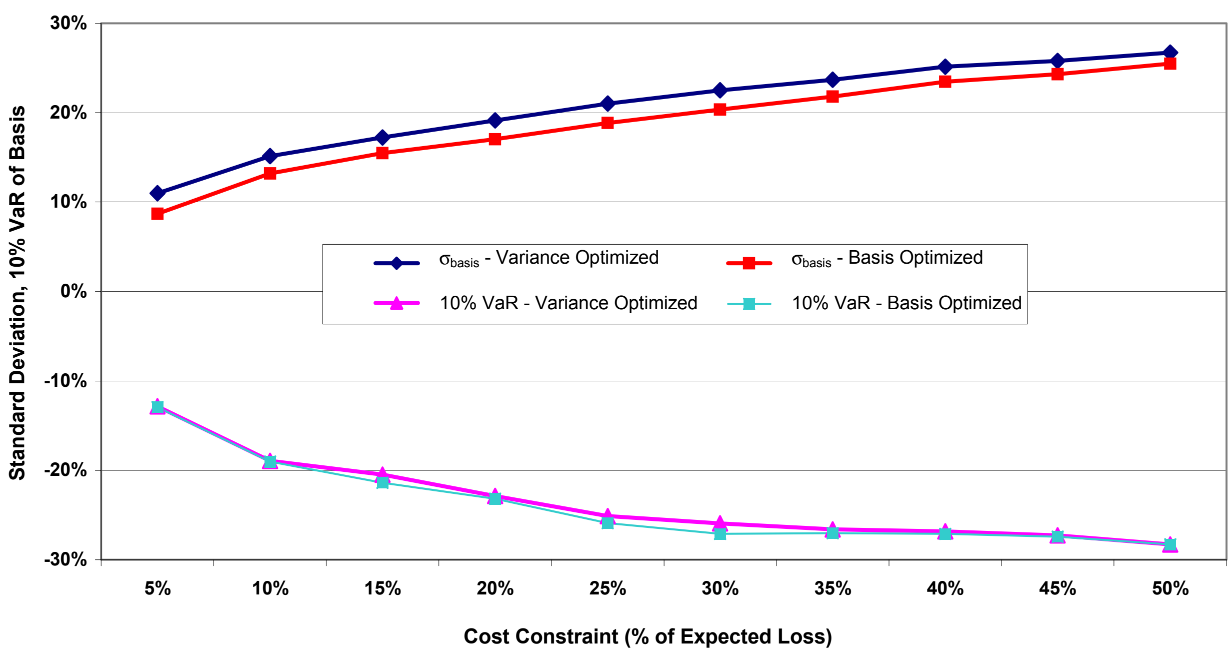

Appendix A

The Applied Insurance Research (AIR) Catastrophe Simulation Models

In this Appendix, we describe AIR's approach to the modeling of natural catastrophes, with a focus on hurricanes. We then discuss how catastrophe modeling technology is used to estimate both index values and individual company loss. A more detailed technical description of the model is available from the authors.

AIR catastrophe models use sophisticated simulation techniques to estimate the probability distribution of losses that result from potential natural catastrophes. A simplified flow chart of the model is shown in Figure A.1. The model first generates the frequency with which events occur, their location and magnitude. After simulated events are generated, they are propagated over the affected area. Local intensity is calculated for every site affected by the event. Next, using detailed information on property locations, values and construction characteristics, the AIR models estimate the probabilities of losses of various sizes. Insured losses are calculated by applying policy conditions to the total damage estimates. This information is then synthesized and further analyzed to assist in risk management.

The AIR Hurricane Model

The hurricane loss estimation methodology employed by AIR is based on well-established scientific theory in meteorology and wind engineering. The simulation models were developed through careful analyses and synthesis of all available historical information and they incorporate statistical descriptions of a large number of variables that define both the originating event (e.g., hurricane) and its effect on structures. The models are validated and calibrated through extensive processes of both internal and external peer review, post-disaster field surveys, detailed client data from actual events and overall reasonability and convergence testing. The AIR hurricane model has been used by the insurance industry since 1987 and is well known for its reliability and the credibility of the loss estimates it generates.

AIR employs Monte Carlo simulation, a well-known statistical technique, to generate simulated storms. Monte Carlo simulation involves an iterative process using, in each simulation, a set of values stochastically drawn from the probability distributions governing each of the random variables being analyzed. In the AIR hurricane model, the random variables being analyzed are landfall location and hurricane frequency, as well as the primary meteorological parameters of each simulated storm (see "Hurricane Event Generation" below). Theoretical probability distributions are fit to the historical data using goodness-of-fit tests and AIR's meteorological expertise. By repeating the simulation process, a sample of more than eighteen thousand storms is generated, each corresponding to a different set of random values assigned to the storm parameters. A sample from a Monte Carlo simulation can be analyzed in ways similar to the ways in which a sample of experimental observations can be analyzed. In particular, a sample from a Monte Carlo simulation can be analyzed statistically to generate probability distributions of losses for individual buildings or portfolios of buildings, given the characteristics of each simulated event.

To estimate the hurricane loss potential, 10,000 annual scenarios of potential hurricane experience were simulated, incorporating over 18,000 simulated events. The first step of the AIR hurricane model is to generate the number of hurricanes estimated to make landfall in the simulated year. The model allows for the possibility of multiple events occurring within a single year. That is, each simulated year may have no, one, or multiple events, just as might be observed in an actual year. For each simulated hurricane, the model first assigns a landfall location and values for each of the modeled meteorological characteristics. It then estimates he potential property damage on the basis of a complete time profile of wind speeds, or windfield, at each location affected by each simulated storm. (The AIR hurricane model also estimates losses from storms that bypass the coast without making actual landfall.)

Data Sources and Analysis

The meteorological sources used to develop the AIR model are databases, information, and publications available from various agencies of the U.S. National Oceanic and Atmospheric Administration (NOAA), including the U.S. National Weather Service (NWS) and the National Hurricane Center. These agencies gather original data on historical hurricanes from such sources as barograph traces from land stations and ships, actual wind records from NWS stations, aircraft reconnaissance flight data, radar data and other pressure and wind reports. These original data are not necessarily consistent. NWS scientists analyze these raw data and use them, along with their professional judgment, to synthesize the primary meteorological characteristics of each historical storm. This final synthesized data are used in developing the AIR model.

AIR then uses statistical estimation techniques to fit various probability distributions to the available meteorological data on historical hurricanes. The distributions employed by the AIR hurricane model are standard statistical distributions that are representative of the underlying historical distributions of the meteorological data. It is not likely, however, that the fitted distributions will duplicate the true underlying distribution of the meteorological data.

Hurricane Event Generation

The first component of the AIR hurricane model provides for the generation of simulated hurricanes. Many thousands of scenario years are generated to produce the complete and stable range of potential annual experience of hurricane activity. For each scenario year, the model generates the fundamental characteristics of each simulated storm, including frequency of occurrence, landfall location and track, and the intensity variables of central pressure, radius of maximum winds and forward speed.

Hurricane Frequency. The model generates the number of hurricanes making landfall for each simulated year from an annual frequency distribution. AIR estimates the parameters of this distribution using the actual hurricane occurrences for the 99 years from 1900 to 1998. The sample includes all landfalling and bypassing hurricanes, where bypassing storms are defined as storms passing sufficiently close to land to cause significant damage.

Landfall Location. Because the values of property exposures vary along the coast, loss estimates can also vary greatly depending on where a hurricane makes landfall. The AIR hurricane model identifies 3,100 landfall points along the coast from Texas to Maine—one for each nautical mile of "smoothed" coastline—and groups these points into sixty-two 50-nautical mile segments of coastline in order to develop a cumulative probability distribution of landfall locations. After tabulating the actual number of historical hurricanes for each 50-nautical mile segment, the actual number of occurrences for each segment is smoothed using a statistical smoothing method used in climatological studies and meteorological judgment. This results in a probability distribution governing landfall location for each segment of modeled coastline.

For illustrative purposes, Figure A.2 shows the number of hurricanes that, since 1900, have made landfall along the Florida coast at each of the twenty 50-nautical mile segments from the Alabama to the Georgia borders. The smoothed frequency distribution ensures that each coastal segment has a non-zero probability of hurricane occurrence (except a few where meteorological or geographical factors prevent hurricanes from making landfall). Therefore, the fact that no hurricane has made landfall at a particular segment in the past does not mean that the AIR hurricane model will simulate no hurricanes for such a segment. Accordingly, the AIR hurricane model allows for the possibility of a hurricane making landfall almost anywhere along the Gulf and Atlantic coasts.

Key Meteorological Characteristics. Once a landfall location is generated for the simulated storm, values are generated for each of the storm's key meteorological characteristics at landfall. For purposes of estimating the probability distributions of these other variables, the coastline from Texas to Maine has been divided into thirty-one 100 nautical mile segments, and each geographic segment has a distinct distribution associated with each variable. Historical storm data corresponding to each of these segments (along with adjacent segments) and each of the variables is fit to theoretical probability distributions. These distributions are used to generate values for each of the simulated storm's key meteorological characteristics, which are:

Central Barometric Pressure. This variable is the lowest sea level barometric pressure at the center of the hurricane. It is the primary determinant of hurricane wind speed. Wind speeds typically increase as the central barometric pressure decreases or, more precisely, as the difference between central pressure and peripheral pressure increases.

Radius of Maximum Winds. The strongest winds in a hurricane are typically found at some distance from the center of the storm. This distance is known as the "radius of maximum winds," and it can range from 5 to over 50 nautical miles. Very intense storms typically have a small radius of maximum winds. A storm making landfall at higher latitudes will typically have a larger radius of maximum winds than one making landfall at lower latitudes.

Forward Speed. This is the rate at which a hurricane moves from point to point. Faster moving storms typically go further inland and are therefore likely to result in losses over a larger area. On the other hand, a faster moving storm will subject any given building to high wind speeds for a shorter duration. In some areas, particularly along the coast, this can lead to lower losses than might otherwise be the case. Both effects are taken into account in the AIR hurricane model.

Storm Track. This is the path the storm takes after landfall, important in determining the properties and structures that are in the path of a hurricane. AIR generates simulated storm tracks based on conditional probability matrices. These allow simulated storm tracks to more closely resemble the curving and recurving tracks that are actually observed.

Local Intensity

Once the model generates the storm characteristics and point of landfall, it propagates the simulated storm along a path characterized by the track direction and forward speed. As the storm moves inland at the forward speed generated as described above, wind speeds begin to diminish due to filling and surface terrain effects. In order to estimate the property losses resulting from the simulated storms, the AIR hurricane model first generates the complete time profile of wind speeds, or windfield, at each location affected by the storm.

Windfield generation requires the following steps:

Maximum Wind Speed. he maximum over-water wind speed is calculated for each simulated hurricane.

Asymmetry Factor. An asymmetry factor, which captures the combined effects of the counter-clockwise motion of hurricane winds and the storm's forward speed, increases wind speeds on the right of the hurricane track, and decreases wind speeds on the left of the track.

Filling Equations. After a hurricane makes landfall, the pressure in the eye of the storm begins to increase, or "fill," causing wind speeds to dissipate. The AIR hurricane model filling equations are a function of geographic region, distance from the coast, and time since landfall. The wind speed at the eye of the storm at any point in time is thus dependent upon the number of hours since landfall.

Adjustment of Wind Speeds for Surface Friction. Each location is assigned an adjustment factor, or friction coefficient, to account for the effects of the local terrain. The horizontal drag force of the earth's surface reduces wind speeds. The addition of obstacles such as buildings will further degrade winds. Friction coefficients are based on digital land use/land cover data.

Estimation of Damages

Once the model estimates peak wind speeds and the time profile of wind speeds for each location, it generates damage estimates for different types of property exposures by combining the exposure information with wind speed information at each location affected by the event.

To estimate building damage and the associated losses, the AIR hurricane model uses damageability relationships, or damage functions. These damageability relationships have been developed by AIR engineers for a large number of different construction and occupancy classes, each designed to provide insight into the wind resistivity of a building.

AIR engineers have developed separate damageability relationships for building contents, with contents damageability a function of the building damage. A third set of functions is used to estimate time element damageability, a function of damage to the building, the time needed to repair or reconstruct the building to usable condition, and the *per diem* expense incurred as a result of the building being unusable or uninhabitable.

Separate damageability relationships for each of building and contents provide estimates of the mean, or expected, damage ratio corresponding to each wind speed as well as probability distributions around such mean. In the case of building damageability, the damage ratio is the dollar loss to the building divided by the corresponding replacement value of the building. For contents, it is the dollar loss to the contents divided by the replacement value of the contents. For time element, the number of calendar days that the building is uninhabitable or unusable is estimated based on the building damage ratio. To calculate business interruption losses, the number of calendar days of effective downtime is multiplied by a *per diem* factor. For both mean damage ratios, the probability distribution of damage ranges from no damage to complete destruction, with probabilities assigned to each level of damage in between. The model estimates non-zero probabilities of zero and one hundred percent loss, as is consistent with empirical observation. A high degree of variability in damage is sometimes observed even within a very small geographic area. AIR damageability relationships attempt to capture this variability.

AIR engineers have developed and refined the damageability relationships over a period of several years. They incorporate documented studies by wind engineers and other experts both within and outside AIR. They also incorporate the results of post-hurricane field surveys performed by AIR engineers and others, and by the analysis of actual loss data provided to AIR by client companies.

Insured Loss Module

In this last component of the catastrophe model, insured losses are calculated by applying the policy conditions to the total damage estimates. Policy conditions may include deductibles by coverage, site-specific or blanket deductibles, coverage limits and sublimits, loss triggers, coinsurance, attachment points and limits for single or multiple location policies, and risk specific reinsurance terms.

Model Output

After all of the insured loss estimations have been completed, they can be analyzed in ways of interest to risk management professionals. For example, the model produces complete probability distributions of losses, also known as exceedence probability curves. Output includes probability distributions of gross and net losses for both annual aggregate and annual occurrence losses. The probabilities can also be expressed as return periods. That is, the loss associated with a return period of 10 years is likely to be exceeded only 10 percent of the time or, on average, in one year out of ten.

Output may be customized to any desired degree of geographical resolution down to location level, as well as by line of business, and within line of business, by construction class, coverage, etc. The model also provides summary reports of exposures, comparisons of exposures and losses by geographical area, and detailed information on potential large losses caused by extreme "tail" events.

Validation and Peer Review of the AIR Models

AIR scientists and engineers validate the models at every stage of development by comparing model results with actual data from historical events. The simulated event characteristics parallel patterns observed in the historical record and resulting loss estimates correspond closely to actual claims data provided by clients. Internal peer review is a standard operating procedure and is conducted by the AIR professional staff of over 50 scientists and engineers, one third of whom hold Ph.D. credentials in their area of expertise. AIR models have also undergone extensive external review, beginning with Dr. Walter Lyons' systematic review of the AIR hurricane model in 1986. Dr. Lyons is an expert meteorologist and consultant with over 24 years of experience and over 130 published book chapters and articles.

Probably the most extensive catastrophe model approval process established to date is that of the Florida Commission on Hurricane Loss Projection Methodology. This Commission was established in 1995 with the mission to "assess the effectiveness of various methodologies that have the potential for improving the accuracy of projecting insured Florida losses resulting from hurricanes and to adopt findings regarding the accuracy or reliability of these methodologies for use in residential rate filings." The Commission has established 40 standards that need to be met before a catastrophe model is acceptable for rate making purposes in the state of Florida. The AIR hurricane model was the only model approved under the 1996 standards, and it has consistently been approved under the standards of subsequent years.

Recent years have witnessed a transfer of catastrophe risk to the capital markets through the issuance of catastrophe, or "cat", bonds. AIR models have been used in the majority of the transactions that have been based on catastrophe modeling. In fact, of the nearly $2 billion of risk capital raised in the last few years, close to 70 percent has been raised in transactions based on AIR catastrophe modeling technology, including modeling of earthquakes, hurricanes, other windstorms. Investors have relied on the research and due diligence performed by the securities rating agencies – Standard & Poor's, Moody's Investors Service, Fitch Investors Service, and Duff & Phelps – to make their investment decisions. As part of the due diligence process, the AIR models and their underlying assumptions undergo extensive scrutiny by outside experts hired by these rating agencies as well as by their own experts. Detailed sensitivity analyses of the major components of the model are performed, stress testing each for model robustness.

Estimating Industry Losses

A fundamental component of AIR analysis is the “industry loss file,” which is a set of estimates of insured industry losses resulting from the events simulated by the AIR catastrophe models. To create the industry loss file, the AIR models estimate the impact of each peril by applying event characteristics to industry-wide exposure data (as opposed to data for a specific insurer). AIR’s estimated property values (see “AlR’s Database of Insured Property Values,” below) for commercial, residential, mobile home and automobile properties are entered into these models and insured losses are estimated. This analysis results in an industry loss file, which consists of the estimated industry losses by county for each of the four business lines, for each simulated event and for each year of simulated events. This industry loss file forms the basis for estimating index values.

For industry loss based indexes, the industry loss file contains the event by event and year by year simulated industry loss values needed to construct both occurrence and aggregate index values. Additionally, the industry loss file contains descriptive information in the form of the simulated parameters such as central pressure, radius of maximum winds and forward speed for each event, which are used in the construction of the parametric indexes studied herein. By running underlying exposure through the model, any index can be simulated. For example the exposures that underlie the GCCI can be quickly analyzed and the index values estimated.

AIR’s Database of Insured Property Values

AIR has developed databases of estimated numbers, types, and values of properties for residential, commercial, mobile home, and automobile insured values in the United States by five-digit ZIP code. These databases have been constructed from a wide range of data sources and reflect the estimated total replacement cost of U.S. property exposures. They are used to estimate total insured property losses. Insured loss estimates are based on assumptions as to the level of deductibles, and how many of the total properties are insured.

The numbers of properties, estimated property values, and other assumptions underlying the database are based on annually updated information. Assumptions specifically regarding insurance policies and trends are based on insurance industry sources including clients, industry organizations, and government studies. The property value databases are developed, maintained and enhanced through an ongoing process of data collection, synthesis and analysis. Much of the information required to develop the estimated values is acquired each year from governmental statistical agencies and private firms that specialize in this type of information. For example, primary data sources in the United States include the U.S. Census Bureau, Dun & Bradstreet, Claritas, the Insurance Information Institute and R.S. Means.

Most data sources supply updated information on an annual basis. While such data sources contain extensive information, AIR has developed internal procedures that select and transform collected data into the required exposure data estimates. These procedures include combining the data from multiple sources and performing appropriate allocations or aggregations of data. For purposes of this analysis, the industry exposure database information is as of July 31, 1998 and no adjustments have been made to reflect the effects of inflation or any other factor since that time.

Estimating Company Losses

For each company in this study, AIR received information on the exposures as described earlier. Where detailed classifications were not provided, AIR assumed industry average characteristics. This exposure information was input into the model described above and, using the same catalogue of events that generated the industry losses, individual company losses were determined. The results are individual company losses, industry loss and event characteristics for each simulated event.

Figure A.1
Flow Chart of the AIR Model

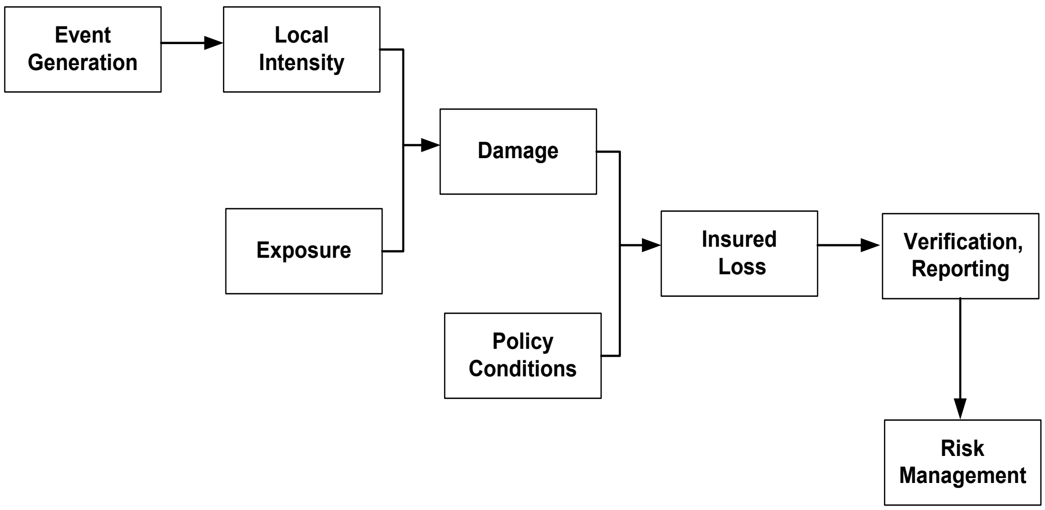

Figure A.2
Number of Hurricanes Making Landfall in Florida: 1900-1998

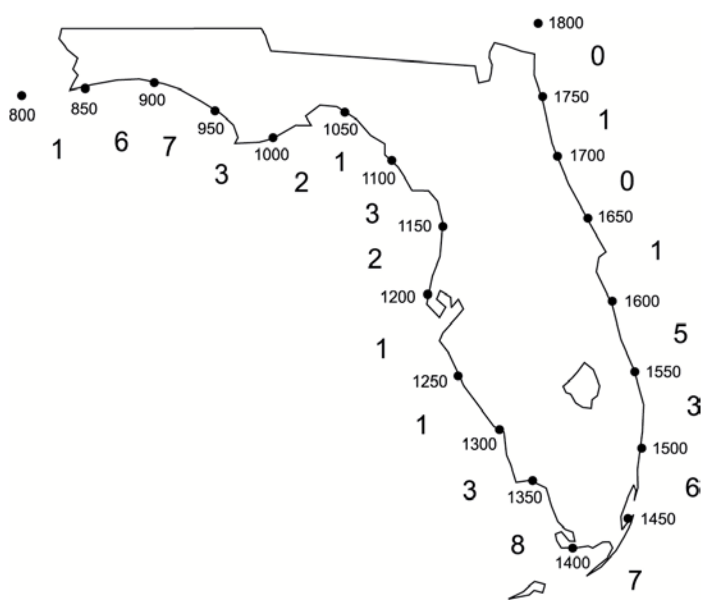

Table A.1
Premium to Expected Payout: Florida CBOT Options and CAT Bonds

A. Florida CBOT Call Spreads

Date	Contract	Premium	Lower Strike	Upper Strike	No. of Contracts	Prem to E[Payout]
Feb-96	Sept/Dec	10,000	80	100	10	6.30
Aug-96	Sept	3,600	40	60	10	1.64
Aug-96	Sept	2,400	40	60	10	1.09
Jul-97	Sept/Dec	69,120	80	100	216	2.01
Jul-97	Sept/Dec	13,600	80	100	40	2.14
Jul-97	Sept/Dec	13,600	80	100	40	2.14
Jul-97	Sept	2,200	100	120	10	2.80
Jul-97	Sept	1,200	100	120	5	3.06
Aug-97	Sept/Dec	8,500	80	100	25	2.14
Sep-97	Sept	1,300	100	120	5	3.31
Dec-97	Dec	600	80	100	30	0.42
Dec-97	Dec	700	80	100	30	0.49
					Average	2.30
Source: Chicago Board of Trade and Applied Insurance Research					**Median**	2.14

B. Catastrophe (CAT) Bond Issues

Date	Transaction Sponsor	Spread Premium	Prob of 1st $ of Loss	E[L | L >0]	Expected Loss	Prem to E[Loss]	Risk
Mar-00	SCOR	2.7%	0.19%	57.89%	0.11%	24.55	Eathquake, Windstorm
Mar-00	SCOR	3.70%	0.29%	79.31%	0.23%	16.09	Eathquake, Windstorm
Mar-00	SCOR	14.00%	5.47%	59.23%	3.24%	4.32	Eathquake, Windstorm
Mar-00	Lehman Re	4.50%	1.13%	64.60%	0.73%	6.16	Earthquake
Nov-99	American Re	2.95%	0.17%	100.00%	0.17%	17.35	Hurricane & Earthquake
Nov-99	American Re	5.40%	0.78%	80.77%	0.63%	8.57	Hurricane & Earthquake
Nov-99	American Re	8.50%	0.17%	100.00%	0.17%	50.00	Hurricane & Earthquake
Nov-99	Gerling	4.50%	1.00%	75.00%	0.75%	6.00	Earthquake
Jun-99	Gerling	5.20%	0.60%	75.00%	0.45%	11.56	Hurricane: Multiple Event
Jun-99	USAA	3.66%	0.76%	57.89%	0.44%	8.32	Single Hurricane
Jul-99	Sorema	4.50%	0.84%	53.57%	0.45%	10.00	Earthquake, Typhoon
Jul-98	Yasuda	3.70%	1.00%	94.00%	0.94%	3.94	Typhoon
Mar-99	Kemper	3.69%	0.58%	86.21%	0.50%	7.38	Earthquake
Mar-99	Kemper	4.50%	0.62%	96.77%	0.60%	7.50	Earthquake
May-99	Oriental Land	3.10%	0.64%	66.04%	0.42%	7.35	Earthquake
Feb-99	St. Paul/ F&G Re	4.00%	1.15%	36.52%	0.42%	9.52	Aggregate Cat
Feb-99	St. Paul/ F&G Re	8.25%	5.25%	54.10%	2.84%	2.90	Aggregate Cat
Dec-98	Centre Solutions	4.17%	1.20%	64.17%	0.77%	5.42	Hurricane: Multiple Event
Dec-98	Allianz	8.22%	6.40%	56.41%	3.61%	2.28	Windstorm and Hail
Aug-98	X.L./MidOcean Re	4.12%	0.61%	63.93%	0.39%	10.56	Cat: Multiple Event
Aug-98	X.L./MidOcean Re	5.90%	1.50%	70.00%	1.05%	5.62	Cat: Multiple Event
Jul-98	St. Paul/ F&G Re	4.44%	1.21%	42.98%	0.52%	8.54	Aggregate Cat
Jul-98	St. Paul/ F&G Re	8.27%	4.40%	59.09%	2.60%	3.18	Aggregate Cat
Jun-98	USAA	4.16%	0.87%	65.52%	0.57%	7.30	Single Hurricane
Mar-98	Centre Solutions	3.67%	1.53%	54.25%	0.83%	4.42	Hurricane: Multiple Event
Dec-97	Tokio Marine & Fire	2.09%	1.02%	34.71%	0.35%	5.90	Earthquake
Dec-97	Tokio Marine & Fire	4.36%	1.02%	68.63%	0.70%	6.23	Earthquake
Jul-97	USAA	5.76%	1.00%	62.00%	0.62%	9.29	Single Hurricane
Aug-97	Swiss Re	2.55%	1.00%	45.60%	0.46%	5.59	Earthquake
Aug-97	Swiss Re	2.80%	1.00%	46.00%	0.46%	6.09	Earthquake
Aug-97	Swiss Re	4.75%	1.00%	76.00%	0.76%	6.25	Earthquake
Aug-97	Swiss Re	6.25%	2.40%	100.00%	2.40%	2.60	Earthquake
					Average	9.09	
Source: Goldman Sachs & Co.					**Median**	6.77	

Appendix B
Counties Composing Each Region in Florida

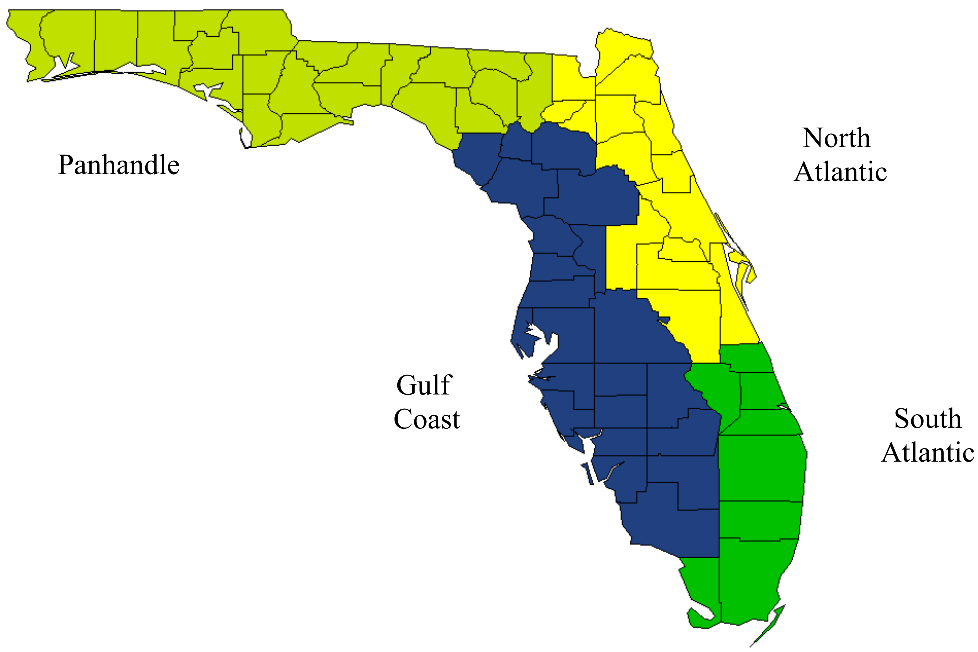

References

Ahn, D., et al., 1999, “Optimal Risk Management Using Options,” *Journal of Finance* 54: 359-375.

American Academy of Actuaries, 1999, “Evaluating the Effectiveness of Index-Based Insurance Derivatives in Hedging Property/Casualty Insurance Transactions,” Report of the Index Securitization Task Force (Washington, DC).

Applied Insurance Research, 1999, *AIR Tropical Cyclone Model: United States Region*, Technical Document TCUS 9904 (Boston).

Artzner, Philippe, Freddy Delbaen, Jean Marc Heber, and David Heath, 1999, “Coherent Measures of Risk,” *Mathematical Finance* 9: 203-228.

Berger, Lawrence, J. David Cummins, and Sharon Tennyson, 1992, "Reinsurance and the Liability Insurance Crisis," *Journal of Risk and Uncertainty* 5: 253-272.

Canter, Michael, Joseph B. Cole, and Richard L. Sandor, 1997, “Insurance Derivatives: A New Asset Class for the Capital Markets and a New Hedging Tool for the Insurance Industry,” *Journal of Applied Corporate Finance,* 10(3): 69-83.

Chookaszian, Jeffrey and Thomas Ward, 1998, “Risk Securitization Products On The Rise,” *National Underwriter (Property & Casualty/Risk & Benefits Management)* 102(20): 9, 23.

Cummins, J. David, Christopher M. Lewis, and Richard D. Phillips, 1999, “Pricing Excess of Loss Reinsurance Against Catastrophic Loss,” in Kenneth A. Froot, ed., *The Financing of Catastrophe Risk* (Chicago: University of Chicago Press).

Cummins, J. David and Mary A. Weiss, 2000, “The Global Market for Reinsurance: Consolidation, Capacity, and Efficiency,”*Brookings-Wharton Papers on Financial Services* (2000), forthcoming.

Doherty, Neil A., 1997, “Financial Innovation in the Management of Catastrophe Risk,” *Journal of Applied Corporate Finance* 10: 84-95.

Dowd, Kevin, 1999, “Financial Risk Management,” *Financial Analysts Journal* 55: 65-71.

Ederington, Louis, 1979, “The Hedging Performance of the New Futures Markets,” *Journal of Finance* 34 (March): 157-170.

Engle, Robert F., and Simone Manganelli, 1999, “CAViaR: Conditional Value at Risk by Quantile Regression,” NBER Working Paper No. W7341, Cambridge, MA.

Froot, Kenneth A., 2001, “The Market for Catastrophe Risk: A Clinical Examination,” *Journal of Financial Economics* 60: 529-571.

Froot, Kenneth A., 1998a, *The Evolving Market for Catastrophic Event Risk* (New York: Marsh and McLennan Securities).

Froot, Kenneth A., 1998b, “Mid Ocean Limited – Trading Catastrophe Index Options,” Harvard Business School Case 9-278-073 (Boston: Harvard Business School Publishing).

Froot Kenneth A. and Paul J. G. O'Connell, 1999, "The Pricing of U.S. Catastrophe Reinsurance," in Kenneth Froot, ed., *The Financing of Catastrophe Risk* (Chicago: The University of Chicago Press).

Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein, 1993, "Risk Management: Coordinating Investment and Financing Policies," *Journal of Finance,* 68: 1629-1658.

Geman, Hélyette, ed., 1999, *Insurance and Weather Derivatives: From Exotic Options to Exotic Underlyings* (London: Risk Books).

Goldberg, D.E., 1989, *Genetic Algorithms In Search, Optimization and Machine Learning* (Reading: Addison-Wesley.

Harrington, Scott E. and Greg Niehaus, 1999, "Basis Risk with PCS Catastrophe Insurance Derivative Contracts," *Journal of Risk and Insurance* 66: 49-82.

Jaffee, Dwight M. and Thomas Russell, 1997, "Catastrophe Insurance, Capital Markets, and Uninsurable Risks," *Journal of Risk and Insurance 64: 205-30.*

Kingdon, J. and K. Feldman, 1995, "Genetic Algorithms and Applications to Finance," *Applied Mathematical Finance* 2: 89-116.

Kunreuther, Howard and Vivek Bantwal, 1999, "A CAT Bond Premium Puzzle," working paper, Wharton Financial Institutions Center, The University of Pennsylvania, Philadelphia.

Litzenberger, R.H., D.R. Beaglehole, and C.E. Reynolds, 1996, "Assessing Catastrophe Reinsurance-linked Securities as a New Asset Class," *Journal of Portfolio Management* (December): 76-86.

Major, John A., 1999, "Index Hedge Performance: Insurer Market Penetration and Basis Risk," in Kenneth A. Froot, ed., *The Financing of Catastrophe Risk* (Chicago: University of Chicago Press).

Merton, Robert and Andre Perold, 1993, "The Theory of Risk Capital in Financial Firms," *Journal of Applied Corporate Finance*.

Pinter, Janos D., 1996, *Global Optimization in Action: Continuous and Lipschitz Optimization – Algorithms, Implementations and Applications* (Norwell, MA: Kluwer Academic Publishers).

Raviv, Arthur, 1979, "The Design of an Optimal Insurance Policy", *American Economic Review* 69: 84-96.

Santomero, Anthony M., 1997, "Commercial Bank Risk Management: An Analysis of the Process," *Journal of Financial Services Research* 12: 83-115.

SwissRe, 1997, "Too Little Reinsurance of Natural Disasters in Many Markets," *Sigma* 7: 3-22.

Varetto, Franco, 1998, "Genetic Algorithms Applications in the Analysis of Insolvency Risk," *Journal of Banking and Finance* 22: 1421-1439.