



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Journal of Financial Economics 71 (2004) 77–111

JOURNAL OF
Financial
ECONOMICS

www.elsevier.com/locate/econbase

The basis risk of catastrophic-loss index securities[☆]

J. David Cummins^{a,*}, David Lalonde^b, Richard D. Phillips^c

^a*The Wharton School, University of Pennsylvania, Philadelphia, PA 19104, USA*

^b*Applied Insurance Research, Boston, MA 02199, USA*

^c*Robinson College of Business, Georgia State University, 35 Broad Street, Atlanta, GA 30303, USA*

Received 15 January 2001; accepted 21 August 2002

Abstract

Using a windstorm simulation model developed by Applied Insurance Research, we analyze the effectiveness of catastrophic-loss index options in hedging hurricane losses for Florida insurers. The results suggest that insurers in the two largest size quartiles can hedge losses almost as effectively using contracts based on four intrastate indices 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 intrastate indices, but most insurers in the smallest quartile would encounter significant basis risk. Hedging using a statewide loss index is effective only for the largest insurers.

© 2003 Elsevier B.V. All rights reserved.

JEL classification: G19; G22; G29; G39

Keywords: Catastrophe risk; Basis risk; Insurance; Hedging; Options

[☆]The authors are grateful to Karen Clark of Applied Insurance Research for her support of this project. We are also indebted to Robert W. Klein for his assistance in obtaining the Florida residential exposure property data used in this study. We are grateful for helpful comments on earlier drafts of this paper from Kenneth Froot, John Major, Stephen D. Smith, and from participants at the 2002 Risk Theory Seminar and the 2001 National Bureau of Economic Research Insurance Project meeting. We thank Andrew Alper and Andrew Kaiser of Goldman Sachs for providing access to the catastrophe bond pricing data. Financial support from the Wharton Financial Institutions Center and the Research Program Council at Georgia State University is gratefully acknowledged.

*Corresponding author. Tel.: +1-215-898-5644; fax: +1-215-898-0310.

E-mail address: cummins@wharton.upenn.edu (J.D. Cummins).

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 a total of \$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 (unpublished data from Applied Insurance Research). 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. (A loss of \$100 billion would equal approximately 30% of the equity capital of the U.S. property-liability insurance industry but would be less than 0.5% of the value of U.S. stock and bond markets.) 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. 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). Moreover, because some have suggested that natural catastrophe losses are “zero-beta” events, CAT-loss securities may 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 asset is not a traded asset or commodity, so that prices are not observed. In this regard, CAT securities are analogous to other new derivatives with “exotic underlyings,” such as weather derivatives (Geman, 1999). 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 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. Industry loss-index contracts are also expected to have lower transactions costs and higher liquidity than insurer-specific contracts because it is easier to standardize contracts and report losses on an index versus having a range of contract specifications and triggering criteria depending upon the characteristics of the issuer as with insurer-specific contracts. Thus, 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. 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 nonfinancial firm, occurring in 1999, covered Tokyo earthquake losses for Oriental Land Company, Ltd. [Froot \(2001\)](#) provides further discussion of the growth of this market.

A primary reason for the lack of success of the CBOT index contracts and the general predominance to date of insurer-specific contracts is the perception among insurers that index-linked CAT securities are subject to unacceptable levels of basis risk ([American Academy of Actuaries, 1999](#)). CBOT contracts were available for nine indices—a national index, five regional indices, and three state indices (for California, Florida, and Texas). Although even the single-state indices were viewed by many insurers as potentially having excessive basis risk, there is virtually no empirical evidence about the basis risk of index-linked securities when used in hedging catastrophe risk by specific insurers. The primary objective of this paper is to conduct a comprehensive analysis of the basis risk and hedging effectiveness of index-linked CAT loss securities. We evaluate a Florida statewide loss index and four intrastate indices to provide evidence on whether insurer concern about basis risk of statewide contracts is justified and the degree to which intrastate indices can be used to reduce basis risk and improve hedging effectiveness. We conduct a simulation analysis of hedging effectiveness for 255 insurers accounting for 93% of the insured residential 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. 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. The AIR hurricane model combines actuarial data, vulnerability relations 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 use the simulated loss

experience to analyze the effectiveness of catastrophic loss hedging strategies for the sample insurers.

The analysis focuses on nonlinear 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 nonlinear 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 that 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 events or to use risk management strategies that are less expensive than purchased hedging instruments. Evidence that hedging is costly is provided by Froot (2001), who shows that both reinsurance and CAT bonds generally trade at significant margins above the expected loss covered by the hedge.

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, although real-world reinsurance and CAT contracts are not in fact “perfect” because they usually include provisions such as coinsurance requirements to control moral hazard, as discussed later. 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 intrastate 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 intrastate regional indices, with size quartiles based on total residential property value exposed to loss in Florida. Many insurers in the third largest size quartile also can hedge effectively using the intrastate 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 nontraded asset, and thus can provide guidance for the securitization of other unconventional financial exposures. Our methodological approach 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 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, whether natural or man-made (such as terrorism e.g., see [Cummins and Lewis, 2003](#)).

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 intrastate 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; and (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 database 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. (Data on the nine omitted insurers were not available from the Florida Insurance Commissioner.) The insurers in our sample account for 93% of the total insured residential property values in the state. Thus, our results can be interpreted as generally representative of the entire insurance industry. Further details about the sample are provided in Section 3.

The simulated catastrophic losses for our sample of insurers are generated using the hurricane model developed by Applied Insurance Research. 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., a 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 is well known for its reliability and the credibility of its loss estimates. Further details

on the model are available from the authors and in *Applied Insurance Research* (1999).

The simulation of property losses in the AIR hurricane model 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 windfield 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 types 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 database. 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 databases. The simulations were also conducted using the AIR zip code database exclusively for a random sample of five companies in order to validate the county aggregation approach. The validation tests indicate 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 database as the exclusive source of information.

2.1. 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 intrastate regional indices. The intrastate indices are based on a subdivision of the state into four segments—the Panhandle, Gulf Coast, North Atlantic, and South Atlantic (see Appendix A). The segment

definitions were provided by Applied Insurance Research based on experience with insurance clients including analyses conducted for the United Services Automobile Association (USAA) CAT bond issues of 1997–1999. We hypothesized that four regions 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. For example, a 1998 attempt to launch zip-code level index contracts failed to generate interest among insurers and investors and is currently dormant (see Chookaszian and Ward, 1998). Our 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 the direct and indirect costs of financial distress and convex tax schedules (see, for example, Froot et al., 1993). 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 (Merton and Perold, 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.

We expect insurers to purchase securitized hedging instruments to perform a function similar to 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 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.

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 nonlinear hedges, where the insurer forms a hedge portfolio consisting of a short position in unhedged catastrophe losses and a long position in call option spreads. The nonlinear 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 et al., 1999).¹

2.2. 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

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

where h_j^P is 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

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

where L_j^S is insurer j 's hedged loss using a statewide industry loss index, h_j^S the hedge ratio for the statewide hedge, $L^S = \sum_j L_j$ the statewide losses for the industry, and M_j^S and U_j^S are the lower and upper strike prices for company j 's statewide call spread. Insurer j 's hedged loss under the intrastate hedge is

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

where L_j^R is company j 's losses under the intrastate regional hedge, L_{jr} the unhedged losses of insurer j in region r , h_j^r the hedge ratio for insurer j in region r , L_r^R the industry-wide losses in region r , M_j^r the lower strike price for company j 's region r call option spread, U_j^r the upper strike price for company j 's region r call spread, and $R =$ the number of regions (four 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:

$$B_j^R = (L_j^P - L_j^R), \quad (4)$$

where B_j^R is 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

¹For purposes of comparison with prior work (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 type of portfolio is evaluated in most of the existing hedging literature (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 nonlinear results because linear hedges are not used in the CAT risk market.

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 statewide 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 statewide hedge is given by

$$\begin{aligned} & \text{Minimize} && G_m[L_j^S \mid L_j \in \{L_j \mid L^S > T\}] \\ & h_j^S, M_j^S, U_j^S && \\ & \text{subject to} && h_j^S[W(L^S, M_j^S) - W(L^S, U_j^S)] \leq C_j, \end{aligned} \quad (5)$$

where C_j is 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 12 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_j^i 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 *industry-wide* loss event such that statewide 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. 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. 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.²

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 nonzero thresholds correspond, respectively, to the 23rd, 14th, and 8th percentiles of the

² 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 on the set of losses resulting from large industry-wide events, and we also conducted the EEV analysis using the industry loss distribution criterion. Use of these alternative hedging approaches had no material effect on the results.

Florida residential property loss severity distribution. We focus on three hedging criteria that 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 following objective function: $G_{1T}(L_j^i | L_j \in L^T) = \sigma_T^2[L_j^i(h_j^i, M_j^i, U_j^i) | L_j \in L^T]$ = 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 $L_j^i(h_j^i, M_j^i, U_j^i)$ is 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 is equal to P for the perfect hedge, S for the statewide industry hedge, and R for the intrastate regional hedge, where the latter is a function of 12 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, $L_j^P \neq L_j$.

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). 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 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 ($0 < \alpha < 1$). Stated more formally, defining insurer j 's net loss distribution function under hedge index i and threshold T as $F_{ijT}[L_j^i(h_j^i, M_j^i, U_j^i) | L_j \in L^T] = F_{ijT}(\cdot)$, VaR is defined as

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

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

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). Recent research suggests that EEV-type measures have desirable properties not possessed by value at risk measures (see, e.g., Artzner et al., 1999). EEV is essentially the value of a call option on L_j^i with strike price equal to a specified loss threshold. More formally, the EEV is defined as

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

where L_V is a loss threshold specified by the decision maker. The EEV criterion function is $G_3(L_j^i) = \text{EEV}_j[L_j^i(h_j^i, M_j^i, U_j^i) | L_j > L_V]$. 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 *hedge effectiveness* as the proportionate reduction in the unhedged 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 statewide index is

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

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.,

$$RHE_{jm}^i = \frac{HE_{jm}^i}{HE_{jm}^P} \quad (9)$$

where $i = S$ is statewide hedge and $i = R$ the 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 index hedge basis given by Eq. (4) provides an alternative measure of hedge efficiency.

2.3. 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).³ The evolutionary genetic algorithm is a global optimization technique designed to efficiently investigate the “entire” feasible set of the parameter space. This differs from the objective of conventional calculus-based optimizers, which seek to refine an initial “guess” about the optimal solution 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).

³See Kingdon and Feldman (1995), Varetto (1998), and Engle and Manganelli (1999) for other applications of genetic algorithms in solving financial problems.

Although global optimization algorithms are superior to conventional methods in extensively exploring the space of possible solutions, they are not necessarily as efficient at refining the solution once the region of the global optimum has been identified (Pinter, 1996; Goldberg, 1989). Consequently, we adopt a multi-stage estimation strategy that combines the genetic algorithm with the calculus-based Newton–Raphson algorithm. Combining global optimization and more conventional algorithms is referred to as a *hybrid optimization algorithm* (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 are 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, e.g., 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 1–3 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 10 cost constraints for a company/objective function combination, we perform a visual inspection of the resulting variance, VaR, or EEV reduction frontier to check for discontinuities and reestimate where necessary. Discontinuities were rare and consisted of cases where the optimization failed to find the global minimum.

To provide further information about the quality of the solutions, we conduct a randomization test using alternative starting values for 20 firms—five firms chosen randomly from each of the size quartiles. Using the variance reduction criterion function, the optimization problem is reestimated 30 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 are randomly selected from intervals with lower bounds equal to 50% of the original starting values and upper bounds equal to 150% of the starting values that would have been used in our original five-step approach. That is, for each of the 30 randomizations for each firm, 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. For example, 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 is 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 30 sets of solutions for all ten cost constraints for each of the 20 firms.

The results of the analysis are robust to the randomization of starting values. The strike prices and hedge ratios for the 30 reestimations for each cost constraint tend to fluctuate in narrow bands around the original estimates. We use the solution value of

the objective function as the primary statistic to evaluate the robustness of the results, i.e., we seek 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 compute the *maximum* percentage reduction in the original hedged net loss standard deviation achieved over the 30 runs for each cost constraint. The maximum reductions are then averaged across the ten cost constraints for each insurer. Averaging the cost-constraint average maximums across insurers, the randomization results would have reduced the statewide net loss standard deviations only slightly (by 0.12% percent on average) and likewise would have reduced the regional standard deviations by only 0.53% on average.

We also evaluate the *maximum* reduction in standard deviation achieved for each company over its 300 randomizations. The maximum reduction based on the statewide hedge is less than 1% for 15 firms, less than 2% for 19 firms, and less than 3% for all 20 firms. For the regional hedge, the maximum reduction is less than 1% for 10 of the 20 firms, less than 3 percent for 15 firms, less than 5% for 18 firms, and between 5% and 8% for the remaining two firms. The objective function value was reduced by 1% or more in only about 5% (309) of the 6,000 trials. 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.

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.

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 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. Thus, the study measures hedging effectiveness for residential property insurance. 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. Having a captive can increase access to hedging mechanisms because captives can participate directly in the reinsurance market and often diversify risk by insuring unrelated parties.

The database 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. 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

Table 1

Summary Statistics: 1998 Florida Insured Residential Property Exposure Database

The table summarizes the residential property exposure data for 255 of 264 insurers operating in Florida in 1998. The residential property category includes insurance on single family dwellings, mobile homes, apartments, and condominiums. Exposure for a property insurance policy is defined as the policy limit, i.e., the maximum amount of loss covered by the policy, measured in dollars. The statistics shown here are calculated across insurers by size quartile where quartiles are based on the total statewide exposures of each insurer. Quartile 1 includes the largest insurers and quartile 4 the smallest. Total statewide exposures is the sum of the total exposures of all insurers in each quartile. Statewide market share is the percentage of total statewide exposure in Florida written by an insurer. Number of counties with exposure is the number of Florida's 67 counties where an insurer has non-zero exposures. Insurers are classified as having ocean exposure if they issue policies in counties with coastal exposure. The county market share coefficient of variation is the standard deviation of the insurer's market share across the 67 counties in Florida divided by the insurer's average county wide market share. The county market share Herfindahl index is the sum of the squares of the percentages of an insurer's total exposures in each county. The data set includes 92.8 percent of the total residential property exposures in Florida. Insurers in quartile 1 account for 94.74% of the total exposures in the sample; and quartiles 2, 3, and 4 account for 4.42%, 0.79% and 0.058% of the total exposures in the sample.

Variable	Size quartile	Average	Std. Dev.	Minimum	Maximum
Total statewide exposures	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	58.34	11.36	15.00	67.00
	2	44.23	14.78	9.00	67.00
	3	29.20	19.14	3.00	67.00
	4	12.48	16.26	1.00	67.00
	All insurers	36.16	23.10	1.00	67.00
% of counties with ocean exposure	1	47.1%	9.2%	25.0%	100.0%
	2	52.7%	8.7%	38.6%	81.8%
	3	60.4%	17.0%	26.5%	100.0%
	4	70.6%	26.3%	0.0%	100.0%
	All insurers	57.6%	18.9%	0.0%	100.0%
% of exposures in ocean counties	1	70.1%	16.7%	23.2%	100.0%
	2	71.4%	18.2%	18.6%	99.7%
	3	70.1%	27.2%	8.7%	100.0%
	4	73.6%	31.8%	0.0%	100.0%
	All insurers	71.3%	24.2%	0.0%	100.0%
County market share coefficient of variation	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	836.0	548.0	244.6	2617.1
	2	1262.4	1160.0	302.5	6485.3
	3	2398.9	2033.4	248.2	8922.0
	4	4479.3	3150.3	345.9	10000.0
	All insurers	2235.4	2417.5	244.6	10000.0

Source: Florida Department of Insurance regulatory filings.

distribution of exposures across the companies in the industry is highly skewed, with the top quartile of insurers accounting for 94.7% 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 might 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, respectively. Larger firms also tend to be more diversified in terms of the coefficient of variation of the market share across counties and in terms of the county market share Herfindahl index, which equals the sum of the squares of the percentages of an insurer's total exposures in each county. These results suggest that larger firms should be better able than smaller insurers to hedge efficiently using index-linked contracts.

3.1. 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 database. Robustness checks based on conducting the optimization using the full 10,000 years of experience for a random sample of ten insurers reveal 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 $L_{jkr,t}$ = hurricane losses for company j , in county k_r , located in intrastate region r , for simulation year t , where $j = 1, \dots, 255$; $k_1 = 1, \dots, K_1; \dots; k_4 = 1, \dots, K_4$; $r = 1, \dots, 4$; $t = 1, \dots, 10,000$ (1,000 for most of the analysis), k_r is the k th 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:

$$\text{The "Perfect" Index} = L_{j..t}^P = \sum_{r=1}^R \sum_{k_r=1}^{K_r} L_{jk_r,t} \quad (10)$$

$$\text{The Regional Indices} = L_{..rt}^R = \sum_{j=1}^N \sum_{k_r=1}^{K_r} L_{jk_r,t} \quad (11)$$

$$\text{The State Index} = L_t^S = L_{...t} = \sum_{j=1}^N \sum_{r=1}^R \sum_{k_r=1}^{K_r} L_{jk_r,t} \quad (12)$$

where N is 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.

3.2. Nonlinear (call spread) hedging

As mentioned earlier, the analysis focuses on nonlinear, “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 Eq. (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% to 50%. 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. Fig. 1 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. (The results based on the \$2.5 billion and \$5 billion thresholds are similar and therefore not shown.) 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 statewide hedge, and the regional hedge. The results confirm that hedging with the regional loss indices is more effective than hedging using the statewide 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% of expected losses reduces the conditional net loss variance by 40% using the statewide hedge, 58% using the regional hedge, and 62% 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 Fig. 2, 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

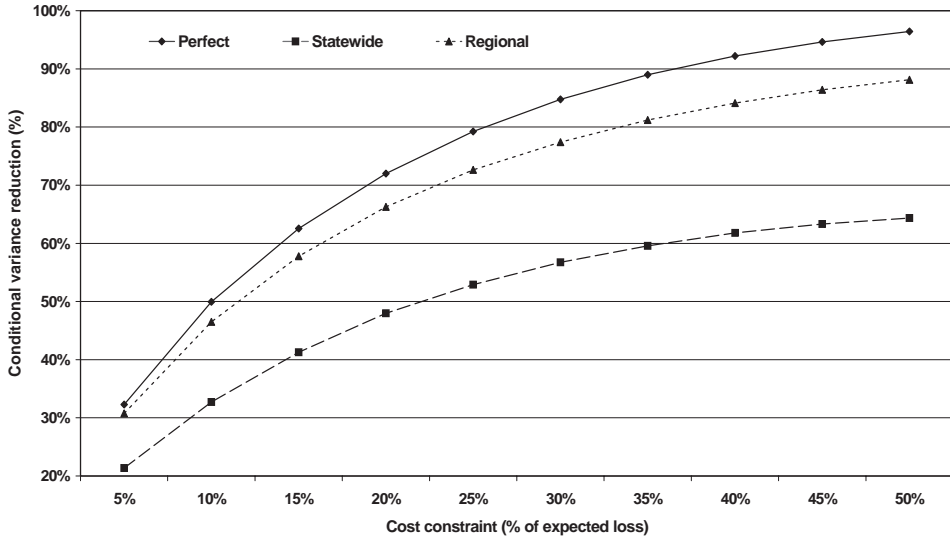


Fig. 1. Conditional variance reduction frontiers: average for insurers in largest size quartile. This figure shows variance-reduction frontiers for the insurers in the largest size quartile, based on exposure to residential property losses in Florida. The frontiers show the percentage reductions in the insurers’ unhedged loss variances using contracts based on three types of loss indices: A “perfect” index based on the insurer’s own losses, regional indices based on total industry losses in each of four Florida regions, and a statewide loss index based on total statewide industry losses. Hedges are constructed assuming that insurers hedge against losses from windstorm events causing total statewide property losses of at least \$1 billion. Variance reduction is shown for each of ten hedging cost constraints, equal to percentages of each insurer’s expected Florida windstorm losses. Points on the frontier are unweighted averages of the percentage variance reduction for the firms in the largest size quartile. The analysis uses 10,000 years of simulated windstorm loss experience from Applied Insurance Research.

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, Fig. 3 shows the frequency distribution of the conditional variance-reduction hedge efficiency in the \$1 billion layer for an expenditure of 15% of expected losses. 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 as in Eqs. (8) and (9). The most striking result based on Fig. 3 is that the regional hedge is at least 95% 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% 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% as effective as the perfect hedge for 36 of the 255 firms, again at the 15% expenditure level with the \$1 billion loss threshold. However, the statewide




Fig. 2. Conditional variance reduction frontiers by insurer size quartile using regional indices. This figure shows variance-reduction frontiers based on unweighted averages for the insurers in each size quartile (quartile 1 contains the largest insurers). The frontiers show the percentage reductions in the insurers' unhedged loss variances using contracts with payoffs triggered by loss indices based on total industry losses in each of four Florida regions. Variance reduction is shown for each of ten hedging cost constraints, equal to percentages of each insurer's expected Florida windstorm losses. Hedges are constructed assuming that insurers hedge against losses from windstorm events causing total statewide property losses of at least \$1 billion. The analysis uses 10,000 years of simulated windstorm loss experience from Applied Insurance Research.

hedge is no more than 50% efficient for 105 of the sample firms, confirming that many firms cannot hedge efficiently based on the statewide index. We also developed a parametric index 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 (results available from the authors).

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 Fig. 4, for an expenditure of 15% 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, Fig. 4 reveals that 70% of the exposures in our sample could be hedged with at least 95% efficiency and 92% could be hedged with at least 90% efficiency using regional index contracts. Even using the less-effective statewide loss index, about 36% of the total

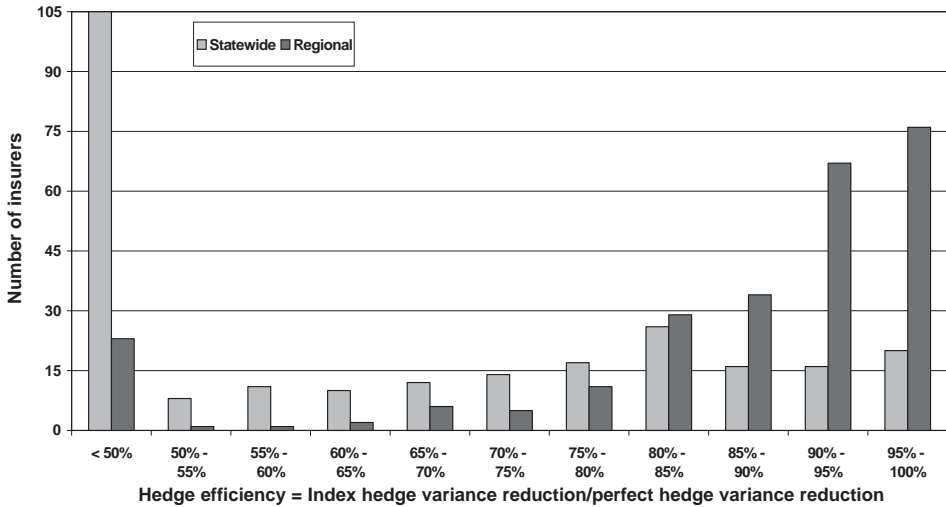


Fig. 3. Conditional variance reduction hedge efficiency by number of firms: hedging cost constraint = 15% of expected loss. This figure shows the variance reduction hedging efficiency of statewide and regional index hedges for 255 property insurers in Florida, tabulated by the number of firms that can hedge with each level of efficiency. Hedge efficiency is the ratio of the variance reduction using index-linked hedges to the variance reduction from the “perfect” hedge based on the insurer’s own losses. The payoff on the statewide hedging contracts is based on the total statewide industry losses from each event, and the payoff on the regional hedging contracts is based on total industry losses in each of four Florida intrastate regions. The results shown are for hedges constructed using a cost constraint of 15% of each insurer’s expected windstorm losses in Florida, assuming that insurers hedge against losses from windstorm events causing total statewide property losses of at least \$1 billion. The analysis uses 10,000 years of simulated windstorm loss experience from Applied Insurance Research.

exposures could be hedged with at least 95% efficiency and 55% could be hedged with at least 90% 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.

We next consider the value at risk (VaR) and expected exceedence value (EEV) hedging criteria. Because the analyses of these two criteria lead to similar conclusions and the EEV has more desirable theoretical properties than the VaR (Artzner et al., 1999), we focus on the EEV. (The VaR results are available from the authors.) Recall that the EEV is the expected loss, conditional on losses exceeding a specified threshold, as in Eq. (7). We report on the EEV based on the 95th percentiles 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 yield similar conclusions. The ratio of the expected CAT loss above the 95th percentile to the total expected CAT loss ranges monotonically from 19% for firms in the first quartile to 31% 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.

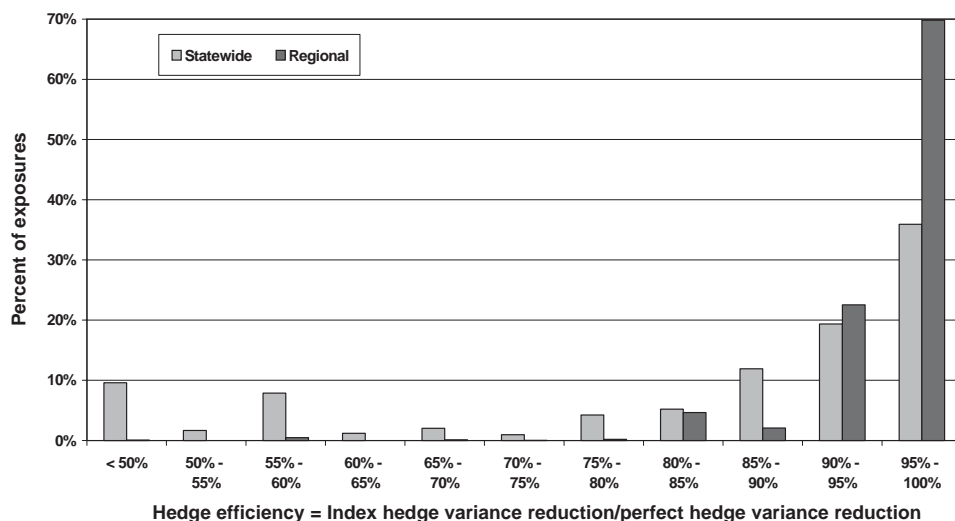


Fig. 4. Conditional variance reduction hedge efficiency by exposures: hedging cost constraint = 15% of expected loss. This figure shows the variance reduction hedging efficiency of statewide and regional index hedges tabulated by the percentages of exposures in Florida that can be hedged with each level of efficiency. The estimates are determined by sorting the 255 insurers in the sample into relative efficiency bins and then, for each sorting bin, summing the total exposures of the insurers in that range and dividing by total statewide exposures in Florida. Hedge efficiency is the ratio of the variance reduction using index-linked hedges to the variance reduction from the “perfect” hedge based on the insurer’s own losses. The payoff on the statewide contracts is based on the total statewide industry losses from an event, and the payoff on the regional contracts is based on total industry losses in each of four Florida intrastate regions. The results are for hedges with a cost constraint of 15% of each insurer’s expected windstorm losses in Florida, assuming that insurers hedge against losses from events causing total statewide losses of at least \$1 billion. The analysis uses 10,000 years of simulated windstorm losses from Applied Insurance Research.

The expected exceedence value (EEV) reduction frontiers for the firms in the largest size quartile are shown in Fig. 5. 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 EEV using the perfect hedge is reduced to zero at an expenditure of approximately 40% 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% of the 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.)

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% of expected losses, the EEV is reduced by 80% using the regional hedge and by 92% using the perfect hedge. Hence, the efficiency of the EEV




Fig. 5. Expected exceedence value reduction frontiers: average for insurers in largest size quartile, EEV threshold = 95th percentile. This figure shows expected exceedence value (EEV) reduction frontiers based on unweighted averages for the insurers in the largest size quartile, with size based on exposure to residential property losses in Florida. EEV is the expected loss conditional on losses exceeding a specified threshold (see Eq. (7)). The figure shows percentage reductions in the insurers’ unhedged EEVs using thresholds equal to the 95th percentile of each insurer’s unhedged loss distribution. EEV reduction is shown for three types of loss indices: a “perfect” index based on the insurer’s own losses, regional indices based on total industry losses in each of four Florida regions, and a statewide index based on total statewide industry losses. EEV reduction is shown for ten hedging cost constraints, equal to percentages of each insurer’s expected Florida windstorm losses. The analysis uses 10,000 years of simulated windstorm loss experience from Applied Insurance Research.

hedges tends to be lower than the efficiency of variance hedges at the higher expenditure levels. Nevertheless, at the 20% expenditure level, 97 of the firms in the sample can hedge with at least 95% EEV efficiency using regional contracts and 138 can hedge with at least 90% efficiency. Also generally comparable to the variance results, 66% of exposures in the sample can be hedged with at least 95% EEV efficiency and 81% can be hedged with at least 90% EEV efficiency, again at the 20% 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 as in Eq. (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 $\sigma_T^2[B_j^i(h_j^i, M_j^i, U_j^i) | L_j \in L^T, L_j^P \neq L_j]$ 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 billion industry-wide loss threshold. The condition $L_j^P \neq L_j$ 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 find that the minimum variance solution often results 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% around zero and, alternatively, plus or minus 5% from zero, where percentages are 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 Fig. 6, which shows the average standard deviation of the basis for the firms in the largest size quartile. 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 $\pm 5\%$ average basis constraint. In each case, regional contracts are used to construct the hedge. 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 smallest size quartile.

Fig. 6 shows that the standard deviation of the basis is about 10% for the 5% cost constraint and then rises to about 25% for the 50% 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% 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%. 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 Fig. 6. First, the basis standard deviations and the absolute values of the 10% VaRs are positively related to the cost

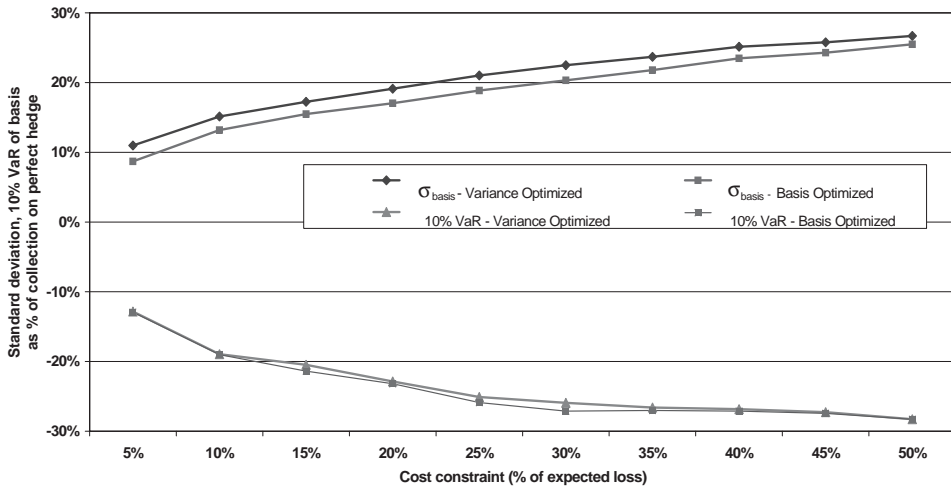


Fig. 6. Standard deviation and 10% VaR of the basis: variance-optimized vs. basis-optimized regional hedges, average for insurers in largest size quartile. This figure shows unweighted averages of the standard deviation of the index-hedge basis and the 10% value at risk (VaR) of the basis, expressed as percentages of the payoffs under the perfect hedge, for insurers in the largest size quartile. Size is based on the insurers' total exposure to residential property losses in Florida. The index hedge basis is the difference between the payout the insurer receives on the optimal regional index hedge contracts and the payout the insurer would receive on the optimal perfect hedge contract, conditional on the perfect hedge contract making a payout. Optimal contracts are determined in two ways: (1) by minimizing the conditional variance of the insurer's net loss distribution (variance optimized); and (2) by minimizing the variance of the index hedge basis (basis optimized). Hedges are constructed to protect against losses from windstorms causing total statewide property losses of at least \$1 billion. Results are shown for ten cost constraints, equal to percentages of each insurer's expected Florida windstorm losses. The analysis uses 10,000 years of simulated windstorm loss experience from Applied Insurance Research.

constraint. The reason for this relation 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% cost constraint, the basis optimized standard deviation is about 13% and the 10% VaR is -20% . This implies that there is a 10% chance of collecting 80% or less of what would have been collected under the perfect hedge for the relatively large losses that are hedged with the 10% cost constraint. For the 20% cost constraint, the 10% VaR is about 25% and the VaR increases (in absolute value) for higher cost constraints.

3.3. Hedge effectiveness and conditional hedging: further discussion

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 90–95% 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 Fig. 6, there is a 10% chance of under-collecting by 15–25% if the cost constraint is in the range that leads to hedging of the most severe losses (cost constraints of 5–20%). 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%, and the USAA CAT bond issues have typically contained a co-payment of 10% or 20%, depending on the issue (Froot, 2001). Because the USAA CAT bond issues have been fully collateralized, these provisions are equivalent to a 100% probability of collecting 80 or 90% of the loss in the call spread coverage layer. This compares to a 90% probability of collecting at least 75–85% with the index options portrayed in Fig. 6.

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 75–90% of a large catastrophic loss in the hedged layer could place the uncovered loss in a range of magnitude that is sustainable without substantial disruption to the firm’s operations, i.e., even a partial recovery could 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%. Although insurers also are likely to face counterparty risk in the market for index-linked CAT securities, properly designed margin and collateral requirements can be used to significantly reduce counterparty risk for exchange-traded contracts. 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.

The preceding analysis has emphasized conditional hedging. That is, we formed hedges on the assumption that insurers hedge large losses, defined as losses arising from an industry-wide event that causes \$1 billion, \$2.5 billion, or \$5 billion in property damage statewide. As a robustness check, we also conduct the analysis on an *unconditional* basis. That is, rather than forming hedges by solving optimization problems where insurers are constrained to minimizing risk above some conditioning point, we solve the hedging problem without imposing the conditioning constraint, potentially allowing for hedges to be established to reduce risk attributable to smaller losses. We then compute the relative efficiency of the optimal hedges based on all four conditioning thresholds (zero, \$1 billion, \$2.5 billion, and \$5 billion) in minimizing the variance of net losses above the \$1 billion industry loss.

The comparison shows 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 leads to the hedging of the largest losses first, followed by the allocation of additional hedging expenditures to progressively smaller losses. We do not observe solutions where funds are 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.

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.

3.4. 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 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. 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 Table 2.

Table 2
Pricing history: CBOT call spreads and CAT bonds

Panel A: Florida Chicago Board of Trade (CBOT) call spreads

Panel A displays all Florida catastrophe call spread contracts traded on the Chicago Board of Trade from 1994–2000. The table shows the month and year of the trade and the time period (months) covered by the contract. The lower and upper strikes for the call spread are expressed in index “points” equal to industry losses divided by \$100 million. For example, an \$8 billion event would have an index value of 80. The settlement value of each point is \$200; a \$9 billion event would lead to a payment of \$2,000 for a contract with lower and upper strikes of 80 and 90, respectively. The contracts are settled on industry Florida loss indices from Property Claims Claims Services. Premium-to-E[Payout] is the ratio of the option premium to the expected payout of the option based on 10,000 years of simulated Florida windstorm losses provided by Applied Insurance Research.

Date	Contract	Option premium (\$)	Lower strike	Upper strike	No. of contracts	Prem-to-E[Payout]
February-96	July–December	10,000	80	100	10	6.30
August-96	July–September	3,600	40	60	10	1.64
August-96	July–September	2,400	40	60	10	1.09
July-97	July–December	69,120	80	100	216	2.01
July-97	July–December	13,600	80	100	40	2.14
July-97	July–December	13,600	80	100	40	2.14
July-97	July–December	2,200	100	120	10	2.80
July-97	July–December	1,200	100	120	5	3.06
August-97	July–December	8,500	80	100	25	2.14
September-97	July–September	1,300	100	120	5	3.31
December-97	October–December	600	80	100	30	0.42
December-97	October–December	700	80	100	30	0.49
					Average	2.30
					Median	2.14

Source: Chicago Board of Trade and Applied Insurance Research

Table 2. (Continued)

Panel B: catastrophe (CAT) bond issues

Panel B displays catastrophe bonds issued from 1997 through March 2000 reported to us by Goldman-Sachs. The spread premium is the annual coupon rate above a reference interest rate (usually one-year LIBOR). The CAT bonds are asset-backed securities where principal payment is released to the issuing insurer on the occurrence of a defined catastrophic loss event. The loss distribution characteristics for each bond are based on simulations conducted by a catastrophe modeling firm. The probability of first \$ of loss is the probability that a contingent payment will be triggered under the bond. The $E[L | L > 0]$ is the expected principal payment to the issuing insurer, conditional on the occurrence of a loss that triggers payment under the bond, expressed as a percentage of the principal of the bond. The expected loss is the product of the probability of first \$ of loss and $(E[L | L > 0])$. Premium to $E[\text{Loss}]$ is the ratio of the spread premium to the expected loss of principal of the bond. The risk is the types of property catastrophes that would trigger payment under the bond.

Date	Transaction sponsor	Spread premium (%)	Prob of first \$ of loss (%)	$E[L L > 0]$ (%)	Expected loss (%)	Prem to $E[\text{Loss}]$	Risk
March-00	SCOR	2.70	0.19	57.89	0.11	24.55	Earthquake, Windstorm
March-00	SCOR	3.70	0.29	79.31	0.23	16.09	Earthquake, Windstorm
March-00	SCOR	14.00	5.47	59.23	3.24	4.32	Earthquake, Windstorm
March-00	Lehman Re	4.50	1.13	64.60	0.73	6.16	Earthquake
November-99	American Re	2.95	0.17	100.00	0.17	17.35	Hurricane & Earthquake
November-99	American Re	5.40	0.78	80.77	0.63	8.57	Hurricane & Earthquake
November-99	American Re	8.50	0.17	100.00	0.17	50.00	Hurricane & Earthquake
November-99	Gerling	4.50	1.00	75.00	0.75	6.00	Earthquake
June-99	Gerling	5.20	0.60	75.00	0.45	11.56	Hurricane: Multiple Event
June-99	USAA	3.66	0.76	57.89	0.44	8.32	Single Hurricane
July-99	Sorema	4.50	0.84	53.57	0.45	10.00	Earthquake, Typhoon
July-98	Yasuda	3.70	1.00	94.00	0.94	3.94	Typhoon
March-99	Kemper	3.69	0.58	86.21	0.50	7.38	Earthquake
March-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
February-99	St. Paul/F&G Re	4.00	1.15	36.52	0.42	9.52	Aggregate Cat
February-99	St. Paul/F&G Re	8.25	5.25	54.10	2.84	2.90	Aggregate Cat
December-98	Centre Solutions	4.17	1.20	64.17	0.77	5.42	Hurricane: Multiple Event
December-98	Allianz	8.22	6.40	56.41	3.61	2.28	Windstorm and Hail
August-98	XL/MidOcean Re	4.12	0.61	63.93	0.39	10.56	Cat: Multiple Event
August-98	XL/MidOcean Re	5.90	1.50	70.00	1.05	5.62	Cat: Multiple Event
July-98	St. Paul/F&G Re	4.44	1.21	42.98	0.52	8.54	Aggregate Cat
July-98	St. Paul/F&G Re	8.27	4.40	59.09	2.60	3.18	Aggregate Cat

June-98	USAA	4.16	0.87	65.52	0.57	7.30	Single Hurricane
March-98	Center Solutions	3.67	1.53	54.25	0.83	4.42	Hurricane: Multiple Event
December-97	Tokio Marine & Fire	2.09	1.02	34.71	0.35	5.90	Earthquake
December-97	Tokio Marine & Fire	4.36	1.02	68.63	0.70	6.23	Earthquake
July-97	USAA	5.76	1.00	62.00	0.62	9.29	Single Hurricane
August-97	Swiss Re	2.55	1.00	45.60	0.46	5.59	Earthquake
August-97	Swiss Re	2.80	1.00	46.00	0.46	6.09	Earthquake
August-97	Swiss Re	4.75	1.00	76.00	0.76	6.25	Earthquake
August-97	Swiss Re	6.25	2.40	100.00	2.40	2.60	Earthquake
					Average	9.09	
					Median	6.77	

Source: Goldman Sachs & Co.

Table 3

Variance reduction efficiencies: market price hedge relative to fair price perfect hedge

The efficiencies shown in the table are ratios of the variance reduction that could be achieved with hedging contracts priced above the actuarially fair level to the variance reduction from hedging with actuarially fair perfect hedge contracts. The price-to-expected loss ratio for the perfect hedge contracts is 6.8, the median CAT bond mark-up ratio from Table 2, and the price-to-expected loss ratio for the state and regional index contracts is 2.1, the median mark-up ratio for CBOT option spreads from Table 2. Hedges are constructed assuming that insurers purchase hedging contracts to protect against losses from catastrophic windstorm events causing total statewide property losses of at least \$1 billion.

Cost constraint (% of expected loss)	Market/actuarial efficiency		
	Perfect hedge	Statewide hedge	Regional hedge
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

The contractual forms in the non-actuarial analysis are identical to those used in the hedging analysis above, with the difference 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 Table 2.

The results of the non-actuarial hedging analysis are shown in Table 3. Because the results under different hedging strategies lead to the same conclusions, only the variance reduction results are shown. Table 3 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. We do not show the results based on a common markup for all securities because the relations 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. The ratios in the table are averages based on a stratified (by size quartile) sample of the firms in our database. The sample consists of three firms chosen randomly from each size quartile.

The results in Table 3 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 is reduced in comparison with

actuarially fair perfect hedge contracts. For example, if expenditures on hedging are constrained to 25% of expected losses, the market-priced perfect hedge reduces the conditional variance by only 32.8% 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% cost constraint, the market price state hedge reduces the conditional variance by 51.8% of the reduction that could be achieved using actuarially priced contracts, and the comparable reduction for the regional hedge is 63.7%. 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 2 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 2. 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.⁴ Consequently, CAT loss securities are more likely to approach actuarial fairness than reinsurance, particularly for mega-CATs that 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 moderate basis risk.

4. Conclusions

The securities market has responded to the dramatic increase in catastrophic 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

⁴The imperfections of the reinsurance market are further discussed in Berger et al. (1992), Froot and O'Connell (1999), Cummins and Weiss (2000), and Froot (2001).

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 based on 10,000 years of simulated hurricane losses from Applied Insurance Research.

Three indices are analyzed—a “perfect” index consisting of the insurer’s own losses, a statewide industry loss index, and four intrastate 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.

The principal finding is that firms in the three largest Florida market-share quartiles can hedge almost as effectively using intrastate index contracts as they can using perfect hedge contracts, when the objective is to minimize either the conditional variance, the VaR, or the EEV. For example, the hedges based on regional index contracts are at least 90% 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% 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—93% of the total insured property value in the state could be hedged with at least 90% efficiency using the regional contracts because of the heavy concentration of property values insured by the largest and most highly diversified firms.

Analysis of the index hedge basis reveals that insurers would face a 10% probability of collecting 12% less than the loss amount for the lowest cost constraint, ranging upward to 30% 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 5–20%, and that buying reinsurance exposes the hedging insurer to the reinsurer’s credit risk. Thus, hedging through index contracts may 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 markups 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, markups 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.

Appendix A

Counties composing each region in Florida



References

- Ahn, D., Boudoukh, J., Richardson, M., Whitelaw, R., 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, P., Delbaen, F., Heber, J., Heath, D., 1999. Coherent measures of risk. *Mathematical Finance* 9, 203–228.
- Berger, L., Cummins, J., Tennyson, S., 1992. Reinsurance and the liability insurance crisis. *Journal of Risk and Uncertainty* 5, 253–272.
- Canter, M., Cole, J., Sandor, R., 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, 69–83.
- Chookaszian, J., Ward, T., 1998. Risk securitization products on the rise. *National Underwriter (Property & Casualty/Risk & Benefits Management)* 102 (20), 9, 23.
- Cummins, J., Weiss, M., 2000. The global market for reinsurance: consolidation, capacity, and efficiency. *Brookings-Wharton Papers on Financial Services*, Vol. 2000, pp. 159–222.
- Cummins, J., Lewis, C., 2003. Catastrophic events, parameter uncertainty, and the breakdown of implicit long-term contracting: The case of terrorism insurance. *Journal of Risk and Uncertainty* 26, 153–178.
- Cummins, J., Lewis, C., Phillips, R., 1999. Pricing excess of loss reinsurance against catastrophic loss. In: Froot, K. (Ed.), *The Financing of Catastrophe Risk*. University of Chicago Press, Chicago.
- Doherty, N., 1997. Financial innovation in the management of catastrophe risk. *Journal of Applied Corporate Finance* 10, 84–95.
- Cummins, J., Lewis, C., Phillips, R., 1999. Pricing excess of loss reinsurance against catastrophic loss. In: Froot, K. (Ed.), *The Financing of Catastrophe Risk*. University of Chicago Press, Chicago.
- Doherty, N., 1997. Financial innovation in the management of catastrophe risk. *Journal of Applied Corporate Finance* 10, 84–95.
- Dowd, K., 1999. Financial risk management. *Financial Analysts Journal* 55, 65–71.
- Ederington, L., 1979. The hedging performance of the new futures markets. *Journal of Finance* 34, 157–170.
- Engle, R., Manganelli, S., 1999. CAViaR: conditional value at risk by quantile regression. National Bureau of Economic Research Working Paper No. W7341. Cambridge, MA.
- Froot, K., 1998a. The evolving market for catastrophic event risk. Marsh and McLennan Securities, New York.
- Froot, K., 1998b. Mid ocean limited-trading catastrophe index options. Harvard Business School Case 9-278-073. Harvard Business School Publishing, Boston.
- Froot, K., 2001. The market for catastrophe risk: a clinical examination. *Journal of Financial Economics* 60, 529–571.
- Froot, K., O'Connell, P., 1999. The pricing of U.S. catastrophe reinsurance. In: Froot, K. (Ed.), *The Financing of Catastrophe Risk*. The University of Chicago Press, Chicago.
- Froot, K., Scharfstein, D., Stein, J., 1993. Risk management: coordinating investment and financing policies. *Journal of Finance* 68, 1629–1658.
- Geman, H. (Ed.), 1999. *Insurance and Weather Derivatives: From Exotic Options to Exotic Underlyings*. Risk Books, London.
- Goldberg, D., 1989. *Genetic Algorithms In Search, Optimization and Machine Learning*. Addison-Wesley, Boston.
- Harrington, S., Niehaus, G., 1999. Basis risk with PCS catastrophe insurance derivative contracts. *Journal of Risk and Insurance* 66, 49–82.
- Jaffee, D., Russell, T., 1997. Catastrophe insurance, capital markets, and uninsurable risks. *Journal of Risk and Insurance* 64, 205–230.
- Kingdon, J., Feldman, K., 1995. Genetic algorithms and applications to finance. *Applied Mathematical Finance* 2, 89–116.
- Litzenberger, R., Beaglehole, D., Reynolds, C., 1996. Assessing catastrophe reinsurance linked securities as a new asset class. *Journal of Portfolio Management* 22, 7686.

- Major, J., 1999. Index hedge performance: insurer market penetration and basis risk. In: Froot, K. (Ed.), *The Financing of Catastrophe Risk*. University of Chicago Press, Chicago.
- Merton, R., Perold, A., 1993. The theory of risk capital in financial firms. *Journal of Applied Corporate Finance* 6, 16–32.
- Pinter, J., 1996. *Global Optimization in Action: Continuous and Lipschitz Optimization—Algorithms, Implementations and Applications*. Kluwer Academic Publishers, Boston.
- Raviv, A., 1979. The design of an optimal insurance policy. *American Economic Review* 69, 84–96.
- Santomero, A., 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, F., 1998. Genetic algorithms applications in the analysis of insolvency risk. *Journal of Banking and Finance* 22, 1421–1439.