MANAGING RISK USING INDEX-LINKED CATASTROPHIC LOSS SECURITIES

By

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

May 25, 2002

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

J. David Cummins
The Wharton School
Phone: 215-898-5644
Fax: 215-898-0310

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

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

MANAGING RISK USING INDEX-LINKED CATASTROPHIC LOSS SECURITIES

1. Introduction

Hurricane Andrew in 1992 and the Northridge earthquake in 1994 resulted in $30 billion in insured property losses and led insurers to increase significantly their estimates of potential losses from property catastrophes. Insurers seeking to hedge catastrophe (CAT) losses following these events quickly learned that reinsurance markets were inadequate to reinsure events in the Andrew-Northridge loss range and would be even less effective in funding the projected "Big One" in Florida or California, which could cause losses exceeding $100 billion (Swiss Re 1997, Froot 1998a).[1] Although substantial new capacity in the form of equity capital flowed into the reinsurance market during the 1990s, it has become clear that reinsurance is not the most efficient way to handle extremely large, infrequent loss events, due to the existence of reinsurance price and availability cycles and other market imperfections (Jaffee and Russell 1997, Cummins and Weiss 2000). Financial innovation in the securities market has provided a better solution to the CAT loss financing problem – insurance-linked bonds, options, and other types of derivatives. The objective of this chapter is to discuss these new financial instruments and demonstrate how insurers and other firms exposed to CAT risk can use these securities in their risk management programs.

The logic behind the securitization solution is compelling – a $100 billion loss would amount to about 30% of the equity capital in the U.S. insurance market and about 75% of the equity capital of the international reinsurance market. However, such a loss would be less than 0.5 of 1 percent of the U.S. stock and bond markets and an even smaller percentage of the value of global securities markets. Insurance-linked securities not only provide hedging benefits to insurers but also permit non-insurers such as industrial firms to bypass the insurance market for some types of risk. Securitization also would help to mitigate the reinsurance underwriting cycle by providing an alternative source of capacity when the reinsurance market experiences its periodic crises. Moreover, because catastrophe losses are "zero-beta" events, CAT-loss

---

[1]On reinsurance market capacity, see Swiss Re (1997), Cummins and Weiss (2000), and Froot (2001). The "Big One" projection is based on unpublished data from Applied Insurance Research, Boston.

securities provide a valuable new source of diversification for investors (Litzenberger, et al. 1996, Canter, et al. 1997). The development of securitized products covering geographical areas world-wide would enable investors to effectively diversify their holdings of CAT risk securities, permitting CAT risk to be hedged at much lower risk loadings than in today's reinsurance markets.

The development of the CAT-risk securities market has been impeded by the lack of a traded underlying asset or commodity, so that prices are not observed. In the absence of a traded underlying, CAT-risk securities have been structured to pay-off on three types of variables – issuer-specific catastrophe loss criteria, insurance-industry catastrophe loss indices, and *parametric* indices based on the physical characteristics of catastrophic events. The choice of a triggering variable involves a trade-off between moral hazard and basis risk (Doherty 1997). Securities based on insurer-specific (or hedger-specific) losses have low 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.[2]

Nearly all CAT risk securities issued to date that pay off on insurer-specific indices have the same mathematical structure as excess of loss (XOL) reinsurance (Cummins, Lewis, and Phillips 1999), and most such products have been tailored to the needs of the issuer. Although these design features enable the contracts to come close to replicating conventional XOL reinsurance products, they also convey certain limitations. In addition to exposing investors to moral hazard, it is difficult to generate a liquid secondary market in such contracts because they are not standardized. Thus, these contracts expose issuers and investors to liquidity risk, raising the required risk premia. Contracts tailored specifically to the issuer also are characterized by higher transactions costs than more standardized instruments.

Because of the potential advantages of exchange traded contracts in reducing liquidity risk and transactions costs, this chapter focuses on hedging strategies using index-linked CAT risk securities. We

[2] In fact, the perception among insurers that CAT index securities are subject to unacceptable levels of basis risk has been identified as the primary obstacle to the more rapid development of the CAT-loss securities market (American Academy of Actuaries 1999).

propose a hedging strategy aimed at minimizing the variance of the insurer's hedged net losses. Such a strategy is consistent with the standard approaches in both the insurance and actuarial literature.[3] Because the contracts we specify are index-linked, we explicitly analyze their basis risk by conducting three case studies. We utilize data on the actual county-level exposures of three insurers writing homeowners insurance in Florida – a national company with a large market share and two regional companies with progressively smaller market shares. We formulate optimal variance hedges for the three case study companies and study the basis risk of the hedges using the hurricane simulation model developed by Applied Insurance Research (AIR), a leading CAT modeling firm.[4] The AIR model is used to simulate 10,000 years of hurricane experience in Florida, and the results of the optimal hedges are analyzed to measure the degree of basis risk.

Using simulated losses from AIR on the same events for virtually the entire Florida market, we construct industry-wide Florida loss indices to provide triggers for the hedges for the case study companies. Two types of indices are specified – a state-wide index and four intra-state regional indices. The findings indicate that the large insurer can hedge effectively using either the statewide or intra-state indices but that the smaller firms must rely on the intra-state indices to achieve acceptable basis risk.[5]

---

[3]Other hedging strategies, including minimizing the value at risk (VaR) are discussed in Cummins, Lalonde, and Phillips (2002). The conclusions regarding the effectiveness of index-linked hedges are the same using the other two criteria.

[4]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.

[5]There have been three previous empirical studies of the basis risk of insurance-linked securities. Harrington and Niehaus (1999) conduct a time series analysis of the correlation between state-specific loss ratios for a sample of insurers and the PCS CAT loss index and find that PCS derivatives would have provided effective hedges for many homeowners insurers. Major (1999) conducts a simulation analysis of insurer CAT losses based on insurer exposures in Florida and finds that hedging with a statewide CAT index is subject to substantial basis risk. Cummins, Lalonde, and Phillips (2002) analyze hedge efficiency for 255 of the 264 insurers writing homeowners insurance and Florida and find that most insurers in the top thee Florida size quartiles could hedge effectively using regional loss indices. Unlike the present study, Cummins, Lalonde, and Phillips (2002) does not focus on conducting case studies of specific insurers.

The remainder of the chapter is organized as follows: Section 2 provides more details on the CAT loss financing problem and discusses the most promising CAT-loss derivative contracts that have been developed to date. Section 3 discusses the variance hedging strategy. Section 4 provides an overview of the AIR hurricane model and discusses the data and the empirical techniques used to specify optimal hedges and analyze basis risk. Section 5 presents the results, and section 6 concludes the chapter.

2. Catastrophic Risk and Securitization

In this section, we discuss the catastrophic loss financing problem and explain the role of securitization in financing catastrophic losses. We then provide more details on insurance-linked securities and briefly discuss insurer characteristics hypothesized to be related to hedging efficiency.

The Catastrophic Loss Financing Problem

Both the frequency and the severity of property losses due to natural catastrophes have increased dramatically in recent years. During the period 1970-1986, the number of catastrophes averaged about 35 per year. Beginning in 1987, however, the number of catastrophes increased sharply, and from 1990-2001, the number of catastrophes exceeded 100 in every year (SwissRe 2002).[6] From 1970-1986, insured losses from natural catastrophes exceeded $5 billion in only one year, and the average catastrophe loss for this period was $2.6 billion. From 1987-2001, however, insured catastrophe losses exceeded $8 billion in all but two years, and catastrophe losses averaged $14.3 billion per year. Since 1986, insurers have paid out $215 billion in natural catastrophe losses. Although the largest loss, Hurricane Andrew, resulted in only $18 billion in insured property-losses, modeling firms are predicting that losses from a major California earthquake or Florida hurricane could exceed $100 billion. These figures do not include insured losses from man-made catastrophes such as the World Trade Center terrorist attack, which at the time of this writing is

[6]These figures are based on the definition of a catastrophe devised by SwissRe. Swiss Re defines losses as catastrophic if they exceed specified dollar valued thresholds, that vary by type of catastrophe. For insured property catastrophes other than marine and aviation, Swiss Re defines a catastrophe for 2001 as an event causing at least $35.1 million in insured property loss (see Swiss Re 2002).

expected to cost $40-70 billion. Financing losses from both natural and man-made catastrophic events is clearly a growing problem affecting both insurance markets and capital markets more generally.

At first glance, it might seem that the international insurance and reinsurance markets could easily fund a major property catastrophe. The amount of equity capital in the U.S. property-liability insurance industry is about $350 billion, and the amount of capital in the international reinsurance market is about $125 billion. However, most of this capital is committed to backing insurer promises to pay the relatively small, frequent losses covered by the vast majority of insurance and reinsurance policies. Insurance markets are much less efficient in financing large, infrequent events. As a result, the percentage of insured property covered by catastrophe reinsurance is inversely related to the size of the event, and only a small fraction of the property exposure base in hazard prone U.S. states is covered by catastrophe reinsurance (Swiss Re 1997, Froot 2001). Thus, the capacity of the international reinsurance market is clearly inadequate to fund major catastrophes (Cummins and Weiss 2000). In addition, reinsurance markets are subject to price and availability cycles, often resulting in price increases and supply restrictions following catastrophic events (Froot 1998a, Froot and O'Connell 1999).

Raising additional equity capital in the insurance industry would not be an efficient solution to the CAT loss financing problem because holding capital in an insurer or reinsurer is costly (Jaffee and Russell 1997). Capital held in insurers is subject to regulatory and agency costs; and tax and accounting rules also penalize insurers for holding capital to cover infrequent (e.g., once in 50-year) events. Informational asymmetries between insurers and capital markets regarding exposure to catastrophic events and the adequacy of loss reserves also increase the cost of holding additional equity. Finally, "excess" capital not currently committed to projects with short or intermediate time horizons is likely to attract corporate raiders.

Securitization has been developed as a more efficient approach to solving the CAT loss financing problem. Although a $100 billion catastrophe would amount to about 30 percent of the equity capital of the U.S. property-liability insurance industry and at least 75 percent of the equity of the international reinsurance industry, a loss of this magnitude would be less than one-half of 1 percent of the value of stocks and bonds traded in U.S. securities markets. Securities markets also are more efficient than insurance markets in reducing information asymmetries and facilitating price-discovery. Finally, because natural catastrophes are zero-beta events, CAT securities provide a valuable new source of diversification for investors, shifting the efficient investment frontier in a favorable direction (Litzenberger, et al. 1996, Canter, et al. 1997).

CAT Options and Bonds

To date, the most important CAT securities in terms of the amount of risk-capital raised have been catastrophic risk (CAT) bonds. CAT option contracts have also been offered both over-the-counter and on the Chicago Board of Trade (CBOT). This section provides the distinguishing features of CAT bonds and options and discusses the advantages and disadvantages of each type of contract.

The first catastrophe insurance derivative contracts were introduced by the CBOT, which began listing catastrophic loss futures contracts in 1992. The contracts intended for trading, however, were not the futures themselves but options on the futures. The contracts settled on the basis of an industry-wide loss index compiled by Insurance Services Office (ISO). The ISO index proved to be unsatisfactory, because it was too highly aggregated and was released infrequently. As a result, the ISO-based contracts were replaced by options that settled on loss indices compiled by Property Claims Services (PCS), an insurance industry statistical agent. Nine indices were available – a national index, five regional indices, and three state indices (for California, Florida, and Texas). The indices were based on PCS estimates of catastrophic property losses in the specified geographical areas during quarterly or annual exposure periods.[7] Although the CBOT options are no longer traded due to low trading volume (mainly due to lack of interest in the options by insurers), they represent an important innovation and are likely to provide the model for exchange traded options that almost certainly

[7]The indices were defined as the total accumulated losses divided by $100 million. E.g. a 20/40 Eastern call spread would be in the money for a catastrophic loss accumulation in the Eastern region of more than $2 billion (20 points). Each index point was worth $200 on settlement so that one 20/40 call would pay a maximum of $4,000 (20 points times $200 per point).

will be developed in future years. Over-the-counter options also have been traded, although these usually settle on insurer-specific loss criteria (see Swiss Re (2001) for further discussion of alternative instruments).

Most CAT bonds issued to date have been based on insurer-specific criteria. The structure of a typical CAT bond is shown in Figure 1. Capital raised by issuing the bond is invested in safe securities such as Treasury bonds, which are held by a single-purpose reinsurer. This arrangement keeps the transaction off the balance sheet of the issuer and insulates investors from credit risk. The bond-issuer holds a call option on the principal in the single-purpose reinsurer with triggering or strike conditions usually expressed in terms of the issuing insurer's losses from a defined catastrophic event. If the defined event occurs, the bond-issuer can withdraw funds from the reinsurer to pay claims, and part or all of the interest and principal payments are forgiven. 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.

The first successful non-exchange traded capital market product that insured against catastrophe risk was an $85 million CAT bond issued by Hannover Re in 1994 (Swiss Re 2001). The first CAT bond issued by a non-financial firm, occurring in 1999, covers earthquake losses in the Tokyo region for Oriental Land Company, Ltd., the owner of Tokyo Disneyland. Approximately 60 tranches of CAT bonds have been issued since 1993, raising more than $4.0 billion in risk capital.

Index-linked CAT options and issuer-specific CAT bonds can be compared and contrasted in terms of their transactions costs, liquidity, basis risk, and exposure to moral hazard.[8] CAT options are superior toCAT bonds in terms of transactions costs. CAT options can be traded inexpensively on an exchange, whereas CAT bond issues are subject to substantially higher transactions costs for legal, investment, auditing, and tax advice. CAT options also have the potential to generate a very liquid market due to their standardization and the anonymity of traders. Although a liquid market in CAT bonds can also be envisioned,

[8]Although standardized, exchange traded CAT bonds can certainly be envisioned, this discussion contrasts options with the issuer-specific bonds that have dominated the CAT securities market to date.

the bonds issued to date have low market liquidity because they are not standardized and not traded on an exchange.

Index-linked CAT options also are superior to issuer-specific CAT bonds in terms of exposure to moral hazard. The existence of a CAT bond may give an insurer the incentive to relax its underwriting and claims settlement standards, leading to higher-than-expected losses. CAT options, on the other hand, are relatively free of moral hazard because they settle on industry-wide losses rather than the losses of a specific insurer.[9] The primary advantage of insurer-specific CAT bonds over index-linked CAT options is that insurer-specific bonds expose the hedger to less basis risk than do the options. The empirical analysis in this chapter is designed to provide information on the degree of basis risk that would be faced by insurers in hedging with index-linked CAT loss securities.

3. Hedging Strategies

In this chapter, we present the results of three company-specific case studies based on actual insurers writing homeowners insurance in Florida. The case studies are designed to illustrate the use of index-linked products in hedging catastrophic risk and to determine the degree of basis risk the insurers would face when hedging with this type of option contract. We analyze hedges based on a statewide Florida loss index and four intra-state regional indices. The region definitions were provided by Applied Insurance Research based on experience with insurance clients, including analyses conducted in conjunction with the United Services Automobile Insurance (USAA) CAT bond issues of 1997-1999. The intra-state indices are based on a subdivision of the state into four segments (the Panhandle, Gulf Coast, North Atlantic, and South Atlantic). Four regions were chosen as a subdivision of the state that we hypothesized would be sufficient to enable insurers to create effective hedges without incurring the high transactions costs and lack of liquidity that

[9]Index-linked options are not totally free of moral hazard problems because large insurers may have the ability to manipulate the index by over-reporting losses to the statistical agent. However, because concentration in insurance markets is relatively low, over-reporting by a large insurer would significantly diluted at the index level, unlike over-reporting on an insurer-specific instrument; and auditing procedures could be implemented to reduce the probability of successful "cheating."

would likely result from a finer subdivision of the state.[10]

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 focus primarily on non-linear hedges, where the insurer forms a hedge portfolio consisting a short position in unhedged catastrophe losses and a long position in call option spreads. This approach was adopted because the call option spread is the dominant contractual form in both the CAT securities and catastrophe XOL reinsurance markets (see Froot 1998b, Cummins, Lewis, and Phillips 1999).

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.[11]

Hedge Portfolios

The hedging strategies considered here are designed to parallel the structure of excess of loss reinsurance. The insurer is assumed to form a <u>hedge portfolio</u> consisting of a short position in its own unhedged catastrophe losses and a long position in call option spreads on its own losses or on a loss index. Defining insurer j's hedged net loss under loss index i as L_j^i, its net loss under the perfect hedge (i = P) is:

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

where L_j^P = insurer j's hedged loss under the perfect hedge, L_j = insurer j's unhedged loss, h_j^P = the hedge

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

[11]In reality, reinsurance hedges are not really "perfect," so this term is used primarily for convenience. Purchasing reinsurance exposes the buyer to the credit risk of the reinsurer as well as supply and price fluctuations due to the reinsurance underwriting cycle. In addition, reinsurance contracts usually have cost-sharing provisions to mitigate moral hazard problems and have geographical and other limitations that mean reinsurance collections less than fully indemnify the ceding company.

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

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

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

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

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

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

Hedging Objective

The insurer is assumed to construct the hedge with the objective of minimizing a function of L_j^i subject to a cost constraint. Defining the objective function for criterion m as $G_m(L_j^i)$, the optimization problem using a state-wide hedge, for example, is given as:

$$\underset{\{h_j^S, M_j^S, U_j^S\}}{\text{Minimize:}} \quad G_m(L_j^S) \tag{4}$$

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

where C_j = the maximum amount available to insurer j to spend on hedging, and $W(L^S,M_j^S)$ and $W(L^S,U_j^S)$ = the prices of call options on industry losses L^S with strike prices M_j^S and U_j^S, respectively. Thus, the insurer optimizes over 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. The optimization problem for the perfect hedge is defined similarly. The optimization problem for the regional hedge is also analogous to expression (4) except that there are twelve decision variables – four hedge ratios and four sets of lower and upper strike prices. By varying C_j, it is possible to generate an efficient frontier based on each optimization criterion and loss index.

The hedging strategy we adopt is the minimization of the variance of the hedged net losses.[12] The variance reduction criterion gives rise to the objective function: $G_1(L_j^i) = \sigma^2[L_j^i(h_j^i,M_j^i,U_j^i)]$ = the variance of the insurer j's loss net of the payoff on the call option spread using loss index i, where i = P for the perfect hedge, S for the statewide industry hedge, and R for the intra-state regional hedge, where the latter is a function of twelve rather than three variables.

Hedge Efficiency

For each loss index i, we define <u>hedge effectiveness</u> 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_j^i . Under the variance criterion function, for example, the hedge effectiveness of the state-wide index for insurer j is:

$$HE_j^S = 1 - \frac{\sigma^2[L_j^S(h_j^S,M_j^S,U_j^S)]}{\sigma_j^2[L_j]} \qquad (5)$$

Another useful indicator of hedge performance is <u>hedge efficiency</u>, defined as the hedging effectiveness of the index hedge relative to that of the perfect hedge, i.e.,

[12]Other strategies, including Value-at-risk (VaR) reduction, have received considerable attention in the literature as a hedging criterion (e.g., Ahn, et al. 1999). For further discussion of other hedging criteria in the context of index-linked CAT risk hedging see Cummins, Lalonde, and Phillips (2002).

$$RHE_j^i = \frac{HE_j^i}{HE_i^P} \tag{6}$$

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

4. Empirical Methodology

This section discusses the catastrophe exposure data base used in our modeling exercises as well as the Applied Insurance Research (AIR) hurricane simulation model. The section concludes with a discussion of the methodology used to solve the optimization problems to obtain the hedge ratios and strike prices.

The Florida Residential Property Exposure Data

The data base for the study consists of county-level data, obtained from the Florida Insurance Commissioner, on insured residential property values for 255 of the 264 insurers writing property coverage in Florida in 1998.[13] The insurers in the sample account for 93 percent of the total insured residential property values in the state. Thus, the results can be interpreted as representative of the entire insurance industry. The exposure data provide the basis for our case studies of three Florida insurers. We also utilize AIR simulations for the same events over the entire data base to construct the statewide and regional loss indices that provide the triggers for the option hedges.

Catastrophe Loss Simulations

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

The hurricane loss estimation methodology employed by AIR is based on well-established scientific

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

theory in meteorology and wind engineering. The simulation models were developed through careful analyses and synthesis of all available historical information and incorporate statistical descriptions of a large number of variables that define both the originating event (e.g., hurricane) and its effect on insured structures. The AIR hurricane model has been used by the insurance industry since 1987 and is well known for its reliability and the credibility of the loss estimates it generates. The AIR model was the first to meet the standards of the Florida Insurance Commission on Hurricane Loss Projection Methodology.

The structure of the simulation model is summarized in Figure 2. The process begins with a Monte Carlo simulation of the number of storms per year for a 10,000 year simulation period, generating more than 18,000 simulated events. The landfall and meteorological characteristics are then simulated for each storm, where the meteorological characteristics include central barometric pressure, radius of maximum winds, forward speed, storm direction, and storm track. Once the model generates the storm characteristics and point of landfall, it propagates the simulated storm along a path characterized by the track direction and forward speed. In order to estimate the property losses resulting from the simulated storms, the AIR model generates the complete time profile of wind speeds, or windfield, at each location affected by the storm.

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

A fundamental component of the model is AIR's insured property data base. AIR has developed databases of estimated numbers, types, and values of properties for residential, commercial, mobile home, and automobile insured values in the United States by five-digit ZIP code. These databases have been constructed from a wide range of data sources and reflect the estimated total replacement cost of U.S. property exposures. In the present study, AIR's zip code level data on insured property values for companies doing business in Florida were used in the simulations and aggregated to the county level using information supplied by the Florida Insurance Department to protect the confidentiality of AIR's data bases.

Estimating Hedge Ratios and Strike Prices

We adopt a two-stage estimation strategy to solve the optimization problems defined in expression (4). In the first stage, we solve the optimization problems using a *differential evolutionary genetic algorithm* (Goldberg 1989). Genetic algorithms provide a robust search procedure to solve difficult non-linear or non-smooth optimization problems (Goldberg 1989, Pinter 1996). The method was adopted because standard optimization algorithms were not adequate to find global optima for many of the companies in the Florida sample. Although global optimization algorithms such as the differential genetic algorithm are superior to conventional methods in extensively exploring the space of possible solutions, they are not necessarily as efficient as conventional methods at finding an optimal solution once the region where the global optimum is located has been identified (Pinter 1996). Accordingly, the second stage of our estimation methodology was to use the solution results from the genetic algorithm as starting values for a calculus-based optimization using the Newton-Raphson algorithm.[14] The findings reported in the paper are based on the second stage results.

Constructing the CAT Loss Indices

As mentioned above, 10,000 years of hurricane experience were simulated for each firm in the Florida sample. The simulations produce the variables L_{jkrt} = hurricane losses for company j, in county k, located in intra-state region r, for simulation year t, where j = 1, . . . , 255, k = 1, . . . , 67, r = 1, . . . , 4, and t = 1, . . . , 10,000. The simulated losses are then used to construct the following loss indices:

---

[14]For further discussion of estimation, see Cummins, Lalonde, and Phillips (2002).

$$\text{The "Perfect" Index} = L^{P}_{j\cdot\cdot t} = \sum_{r=1}^{R} \sum_{k=1}^{K_r} L_{jkrt} \tag{7}$$

$$\text{The Regional Indices} = L^{R}_{\cdot\cdot rt} = \sum_{j=1}^{N} \sum_{k=1}^{K_r} L_{jkrt} \tag{8}$$

$$\text{The State Index} = L^{S}_{t} = L_{\cdot\cdot\cdot t} = \sum_{r=1}^{R} \sum_{k=1}^{K_r} \sum_{j=1}^{N} L_{jkRt} \tag{9}$$

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

5. Case Studies: Hedging the Risk of Florida Hurricane Losses

Three Florida homeowners insurers were selected for the case studies of index-hedge effectiveness. This section discusses factors expected to be related to hedge effectivnes, describes the characteristics of the companies, and then reports on their ability to hedge catastrophic risk using index-linked contracts.

Factors Related to Hedge Effectiveness

The first factor often mentioned as a determinant of the ability of an insurer to hedge using index-linked products is firm size, either in terms of raw exposures or market share. The hypothesis is that firms with high state-wide market shares need to have exposures in nearly all areas of the state in order to attain a large scale of operations. Such firms are expected to have a significant impact on the industry-wide loss indices and, therefore, relatively low basis risk. A second factor often linked with hedge effectiveness is diversification across the state. The argument is that it is not size per se but rather diversification that determines a firm's ability to hedge using index-linked contracts. A more subtle version of this argument is that firms need good statewide diversification to hedge with state loss indices but only need to have adequate diversification within one or more of the intra-state regions in order to hedge effectively with the regional loss indices. The analysis presented below provides information on the relative importance of size and diversification in determining hedge effectiveness.

The Case Study Companies

The case studies focus on three insurers writing homeowners insurance in Florida. Company A is a large national personal lines insurer that is near the top of the first (largest) size quartile in terms of the total property value exposed to loss in Florida. Company B is a Florida domiciled insurers near the bottom of the first size quartile that advertises itself as being "committed to Florida home and business owners." This company is included because it is relatively small for a quartile 1 firm but is a firm that targets Florida homeowners insurance as one of its core lines of business. Company C is a much smaller firm, near the bottom of the second largest size quartile in terms of homeowners property value exposed to loss in Florida. This company is included to determine whether smaller insurers that have better than average diversification for their size classes are able to hedge effectively using index-linked contracts. Keeping in mind that the firms in the two largest size quartiles account for about 99 percent of the total exposures in Florida, the firms in the study should provide an indication of the possible hedging effectiveness of firms representing a very high proportion of the total property values exposed to loss in the state.

Summary statistics on the case study companies relative to averages for the four Florida size quartiles are shown in Table 1. The quartiles are based on the value of property exposed to loss in Florida, with quartile 1 being the largest and quartile 4 the smallest. Table 5 shows the approximate percentages of total state exposures written by each of the case study companies. Company A is in the 5 percent market-share range, company B is in the 1 percent range, and company C is much smaller, writing less that 0.1 of 1 percent of total state exposures. Thus, the case study firms are representative in size of firms in all size ranges except the lowest quartile of insurers.

The next two columns of Table 1 show that the case study companies are generally more diversified than the average firms in their respective size groups. These two columns show the market share coefficients of variation and Herfindahl indices. Both statistics are based on the distribution of an insurer's percentage exposures by county. The coefficient of variation is the ratio of the standard deviation of the insurer's exposures across the Florida counties divided by the average percentage exposure across counties. Thus, a low coefficient of variation means that the insurer is highly diversified. The market share Herfindahl index is the sum of the squares of the insurer's county market share percentages. An index of 10,000 would mean that the company had all of its business in one county, and smaller values of the Herfindahl index imply higher levels of diversification. By both measures, companies A, B, and C have superior diversification relative to the average insurer in the respective size quartiles (quartile 1 for companies A and B and quartile 2 for company C). These companies also have non-zero exposures in a higher proportion of Florida counties than the average firm in their respective size quartiles. The only instance in which a case study company is less diversified than average is for company B, when the measure of diversification is the proportion of business written in ocean front counties – company B has 92.1 percent of its exposures in ocean front counties compared to 70.1 percent for the average firm in quartile 1. Thus, the case study companies are representative of insurers that are likely to be able to hedge more effectively than average rather than representing a random sample of the companies doing business in the state. We comment briefly on how other firms could improve their hedging effectiveness after presenting the results for the case study firms.

Further information on the exposure distribution of the case study companies is provided in Figures 3 through 5, which show the distribution of company exposures and county market shares of companies A through C, respectively. Company A has a substantial market share in all but a few counties. Yet it is clear that its distribution of exposures is concentrated in a few counties, many of them with ocean front exposure. Nevertheless, the fact that company A has significant market penetration throughout the state suggests that it may be able to hedge effectively, even with state-wide hedging instruments. Company B, by contrast, has significant market shares in only about five counties but has at least some of its exposures in nearly every county in the state. Its heaviest concentration of exposures is located along the Gulf Coast, which generally experiences less costly catastrophes than the South-Atlantic region. This company has exposure characteristics that may lend themselves to effective hedging using regional contracts. Company C is in a similar position. It has significant market penetration in about 6 counties and has no exposure in 14 counties. Likewise, its distribution of exposures is not significantly diversified across that state, being heavily concentrated in several central-Florida counties. As for company B, the prediction based on the market share and exposure maps is that company C probably cannot hedge very effectively with statewide contracts but may be able to hedge effectively using regional contracts.

Optimal Hedges and Hedging Effectiveness

As discussed above, we solved the variance minimization problems in expression (4) for each of the case study companies, using options triggered by both the state-wide and the regional indices (as alternatives). The standard of comparison for the index-linked contracts is the *"perfect hedge*," which pays 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.

Prior to discussing hedging effectiveness, it is useful to provide some information on the characteristics of the hedges. Recall that the variance minimization problems are solved subject to cost constraints, ranging in 5 percent intervals based on the insurers' expected losses. The problems were initially solved based on the assumption that each insurer spends 5 percent of its expected losses on hedging, and then the cost constraint was relaxed by 5 percent and the problem solved again, up to a maximum of 50 percent of expected losses spent on hedging. It is probably unrealistic to assume that an insurer would spend as much as 50 percent of its expected losses on hedging, but these results were included to show a wide range of potential hedging opportunities.

As an example of hedging strategies, we consider insurer A's optimal hedge positions using the regional contracts at the 20 percent cost constraint level. This insurer had no hedging expenditures allocated to reducing catastrophe risk in the Florida Panhandle, a result that was consistent for all levels of the cost constraint. This reflects the fact that company A has relatively low exposure to loss in the Panhandle (see Figure 3) and also that this region of Florida experiences relatively low losses from hurricanes in comparison with the other three regions. The hedging expenditures were allocated 15.8 percent to the North Atlantic region, 21.9 percent to the Gulf Coast region, and 62.3 percent to the South Atlantic region. This again is consistent with the insurer's relative exposures across the regions (its highest exposure is in the South Atlantic) as well as the relative hurricane risk in the regions.

Insurer A's lower strike price in the South Atlantic region for the 20 percent cost constraint was equivalent to an industry-wide loss of $9.4 billion, indicating that most of its hedging expenditures are directed at hedging large loss events in the region where its property exposure base is highest and where the hurricane risk is also highest. The upper strike price is close to the maximum probable industry loss for the region, indicating that the solution is "virtually full hedging over a large deductible." This outcome is consistent with the economic theory of insurance, which finds that hedging large losses with non-linear contracts has the greatest impact in terms of reducing the firm's risk (e.g., Froot 2001).

Insurer A's hedge ratio in the South Atlantic region is 3.2 percent. This implies that it would receive 3.2 percent of the difference between the industry-wide loss in the South Atlantic and the lower strike price for that region, if the South Atlantic industry loss is above the strike. For example, for an industry-wide loss of $12 billion, company A would receive about $88.5 million (recall that the strike price for this contract is $9.4 billion). The strike prices are generally less and the hedge ratios somewhat larger for company A in the North Atlantic and Gulf Coast regions, reflecting the lower probabilities of large events in these regions. That is, there is a trade-off between the event threshold and event probability such that the marginal reduction in the firm's variance per dollar spent on hedging implies lower strike prices in regions where the chance of extreme events is relatively low.

The striking prices and hedge ratios change in predictable ways as the cost constraints are gradually relaxed. For example, in each region, the lower strike price declines monotonically as the cost constraint increases. The hedge ratio does not vary systematically as the cost constraint declines, but generally fluctuates within a relatively narrow band. For example, in the South Atlantic region, company A's hedge ratio is between 2.9 and 3.3 percent for nine of the ten cost constraints (it is 4.4. percent for the lowest cost constraint. The upper strike price also remains relatively constant as a function of the cost constraint, generally at a level in the neighborhood of the maximum probable loss in a specific region. Thus, the hedges adhere closely to the economically optimal ideal of "full insurance above a deductible," with the deductible (strike price) being the primary instrument used by the solution algorithm in finding the optimal hedge.

Hedging effectiveness is shown in Figures 6 through 8, for companies A, B, and C, respectively. The figures show variance reduction frontiers which plot the percentage reduction in the firms' net loss variance (vertical axis) as a function of the amount spent on hedging (horizontal axis). Three frontiers are shown on each chart, for contracts paying off on the perfect hedge (the firm's own losses), the statewide loss index, and the regional loss indices. The average efficiencies (the variance reduction with the index hedges ratioed to the variance reduction with the perfect hedges) are more or less constant across cost constraints. Consequently, we show the average efficiency across all ten cost constraints in a display box on each chart.

Focusing first on company A, we see that there are diminishing marginal returns to hedging – the variance reduction frontiers are increasing and concave in hedge expenditures. It is also clear that company A can hedge effectively using either the statewide or regional indices. For example, at the 20 percent cost constraint, insurer A can reduce its variance by 72.0 percent using the perfect hedge, 68.2 percent using the regional hedge, and 67.0 percent using the statewide hedge. The distance between the hedging curves increases with the amounts spent on hedging. At the 35 percent expenditure level, variance is reduced by 85.4 percent using the perfect hedge, 80.5 percent using the regional index contracts, and 78.4 percent using the statewide index contracts. The average efficiency statistics show that the statewide index contracts would reduce company A's variance by 94.9 percent in comparison with the perfect hedge, and the comparable efficiency statistic for the statewide hedge in 92.9 percent.

Whether the reductions in variance based on the index contracts are in some sense "sufficient" to motivate the use of index-linked contracts rather than company specific contracts is a value judgment, absent any additional information on the market value of hedging using alternative objectives. However, although insurers are accustomed to thinking of hedges as "perfect" in the sense of reinsurance, which has low basis risk, the reduction rather than elimination of risk is typically the objective in more general hedging contexts. Thus, if transactions costs and pricing were sufficiently attractive in the market for insurance-linked securities, hedging with 95 percent efficiency might ultimately be viewed as an attractive alternative to hedging with 100 percent efficiency through more expensive reinsurance contracts, especially keeping in mind that reinsurance is not really perfect due to credit risk, reinsurance cycles, and other factors.

Figures 7 and 8 show the variance-reduction frontiers for companies B and C. The figures show that neither of these companies can hedge effectively using statewide index contracts. Even with 50 percent of expected losses devoted to hedging, the variance reduction is less than 45 percent for both firms, in comparison to more than 90 percent for the perfect hedge contracts. The hedging efficiencies of the state index hedges in comparison with the perfect hedge are only 40.2 and 24.3 percent, respectively, for companies B and C. These results are not surprising given the exposure and market share mappings for the two firms shown in Figures 4 and 5 and provide an indication that concern about the basis risk of the CBOT option contracts was well founded, at least for some insurers.

Although statewide contracts are not a solution to the CAT risk problem for companies B and C, Figures 7 and 8 also show that regional hedging contracts are likely to be viable for these firms. This is particularly noteworthy for company B, where the regional contracts are 97.5 percent as efficient on average as the perfect hedge contracts. The average efficiency for company C also is quite respectable, 93 percent. These results are reasonably predictable given the exposure mappings shown in Figures 4 and 5.

The results suggest two important conclusions about smaller firms in Florida: (1) Regional index-linked contracts may be a viable hedging alternative for many smaller firms, suggesting the possibility that a liquid market in regional Florida contracts may be foreseeable; and (2) smaller firms can manage risk more effectively if they pay more attention to their exposure diversification across the state. A significant proportion of the smaller Florida homeowners insurers are less diversified than insurers B and C and hence could improve their performance and the efficiency of potential hedging solutions by better exposure management, either through underwriting strategies or through pooling arrangements with other small firms.

The results also have implications regarding the relative importance of size and exposure diversification as determinants of hedging effectiveness. The results with company A suggests that size plays a role in enabling firms to hedge effectively using the statewide contract, because it is important to have significant market penetration in all regions of the state in order for the firm's losses to be highly correlated with the industry-wide index. A smaller firm possibly could achieve reasonably high correlation with the statewide index through careful diversification but such firms are much less important in determining the value of the index and thus will always be at a disadvantage. The results also suggest that diversification within regions is sufficient to hedge effectively with the regional indices, i.e., size conveys less of an advantage when the regional contracts are considered.

6. Conclusions

The securities market has responded to the dramatic increase in catastrophe losses over the last fifteen to twenty years by developing innovative new derivative securities. The introduction of insurance-linked securities also has been driven by the increasing recognition that conventional insurance and reinsurance markets do not provide efficient mechanisms for financing losses from low frequency, high severity events. The two most promising types of CAT securities are CAT bonds and options. CAT securities can be designed to pay off on three principal type of triggers – company specific losses, industry-wide aggregate loss indices, and parametric indices relating to a physical measure of the severity of the catastrophe. Most CAT bonds issued to date have been based on issuer-specific triggering criteria, whereas the CBOT option contracts and a few bond issues have been based on industry-wide loss triggers.

Index-linked CAT risk securities are superior to issuer-specific contracts because they are more easily standardized, providing the potential for a much more liquid market than is likely to exist for issuer-tailored contracts. In addition, exchange traded contracts have lower transactions costs than over-the-counter CAT securities and have lower counter-party risk due to the clearance of transactions through the exchange. The principal advantage of issuer-specific contracts over index-linked contracts is that the former have lower basis risk. This chapter illustrates the use of index-linked contracts in hedging catastrophic risk, and models the basis risk of the contracts using actual exposure data of three insurers in the Florida homeowners market.

In our hedging analysis, we form portfolios consisting of a short position in the insurer's unhedged losses and a long position in call option spreads on loss indices. Three indices are analyzed – a "perfect" index consisting of the insurer's own losses, a statewide industry loss index, and four intra-state regional industry loss indices obtained by dividing the state into four quadrants. The hedging criterion analyzed is the minimization of the variance of the insurer's hedged net losses. We measure hedging effectiveness by comparing hedges based on the statewide and intra-state indices with perfect hedges based on each insurer's own losses and define *hedge efficiency* as the ratio of the risk reduction obtained using industry loss index options to the risk reduction obtained using the perfect index.

The principal findings of the study are: (1) Florida insurers can hedge effectively using contracts based on regional loss indices. Four Florida regions are used in the analysis, indicating that as few as four contracts have the potential to generate a liquid market in index-linked CAT securities. The results thus suggest that it is not necessary to have contracts based on smaller geographical areas such as zip codes in order for insurers to construct effective hedges.[15] (2) Hedging through contracts based on statewide indices

[15]The failed Bermuda Insurance Exchange unsuccessfully attempted to promote hedging contracts defined in terms of zip code level loss indices. For further discussion, see Chookaszian and Ward (1998).

is likely to be effective only for relatively large insurers and/or well-diversified smaller insurers in the top two Florida size quartiles. The majority of firms will require sub-state contracts in order to hedge effectively. Nevertheless, the results suggest that a high proportion of the total exposures in Florida could be hedged effectively using index-linked contracts due to the high skewness of the size distribution of Florida insurers.

Three insurers were analyzed, with market shares of approximately 5 percent, 1 percent, and 0.1 of 1 percent, respectively. The firms chosen for the analysis are each relatively well-diversified in comparison to the average insurers in their respective size quadrants. Hence, the results show that even relatively small insurers can hedge effectively using regional contracts, provided that their exposures are reasonably diversified within the Florida regions. An important implication of this findings is that many Florida insurers could reduce their exposure to catastrophic risk and put themselves in a better position to hedge using index-linked derivative contracts if they were to practice better exposure management either through their underwriting decisions or through pooling arrangements with other small insurers.

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 global reinsurance market in dealing with this type of loss. In order for a liquid market in index-linked derivatives to develop, however, a number of obstacles would have to be overcome. These include the following: (1) Better indices of catastrophe losses need to be developed. The development of index-linked catastrophic loss securities has so far been significantly impeded by the lack of suitable indices. Development and maintenance of CAT loss indices might be an appropriate facilitating role for the Federal government to play, by analogy with the provision of temperature readings by the National Weather Service. This would involve setting up electronic reporting services whereby insurers could rapidly and inexpensively report their catastrophe loss experience to a central source which would release damage estimates on intra-state regional losses within every major state subject to catastrophe exposures.

(2) A second factor needed for the development of the index-linked market is the education of

insurance company management. Managers are accustomed to thinking in terms of reinsurance contracts, which have limited basis risk but also are characterized by low liquidity and exposure to the vagaries of the reinsurance underwriting cycle. A change in institutional thinking of the type that has taken place over the past two decades in terms of corporate finance will be needed to acclimate insurance company managers to more sophisticated risk management strategies.

(3) The final important factor that will be mentioned here is the need for development of better pricing models for CAT loss derivatives. The existence of the Black-Scholes model was one of the primary driving forces behind the rapid development of options markets during the 1970s, and models have already been proposed that extend conventional option pricing models to price other new securities. Pricing CAT options is more difficult, however, precisely because these options protect against low frequency, high severity events. These event characteristics imply that "jump risk" is more important in pricing CAT options than in pricing most other types of derivatives and also that the experience data available to test and calibrate the models is very sparse. The lack of a widely-accepted pricing model also may be partly responsible for the high spread premia characterizing most CAT risk contracts issued to date. Thus, the development of a pricing model that would solve the CAT security pricing problem would constitute an invaluable contribution to both the academic and practical finance literatures.

Table 1
Summary Statistics: Case Study Companies versus Quartile Averages

Company	Quartile	% State Exposures*	County MS CoV	County MS Herfindahl	% Expos in Ocean Front Counties	No. of Counties w/ Expo
Company A	1	5.00%	0.363	434	71.4%	67
Company B	1	1.00%	1.320	895	92.1%	65
Company C	2	0.10%	1.759	876	43.6%	53
Quartile 1		94.7%	1.36	836	70.1%	58.3
Quartile 2		4.4%	2.20	1,262	71.4%	44.2
Quartile 3		0.8%	3.35	2,399	70.1%	29.2
Quartile 4		0.1%	5.38	4,479	73.6%	12.5

*Approximate percentages.

Figure 1
The Structure of a CAT Bond

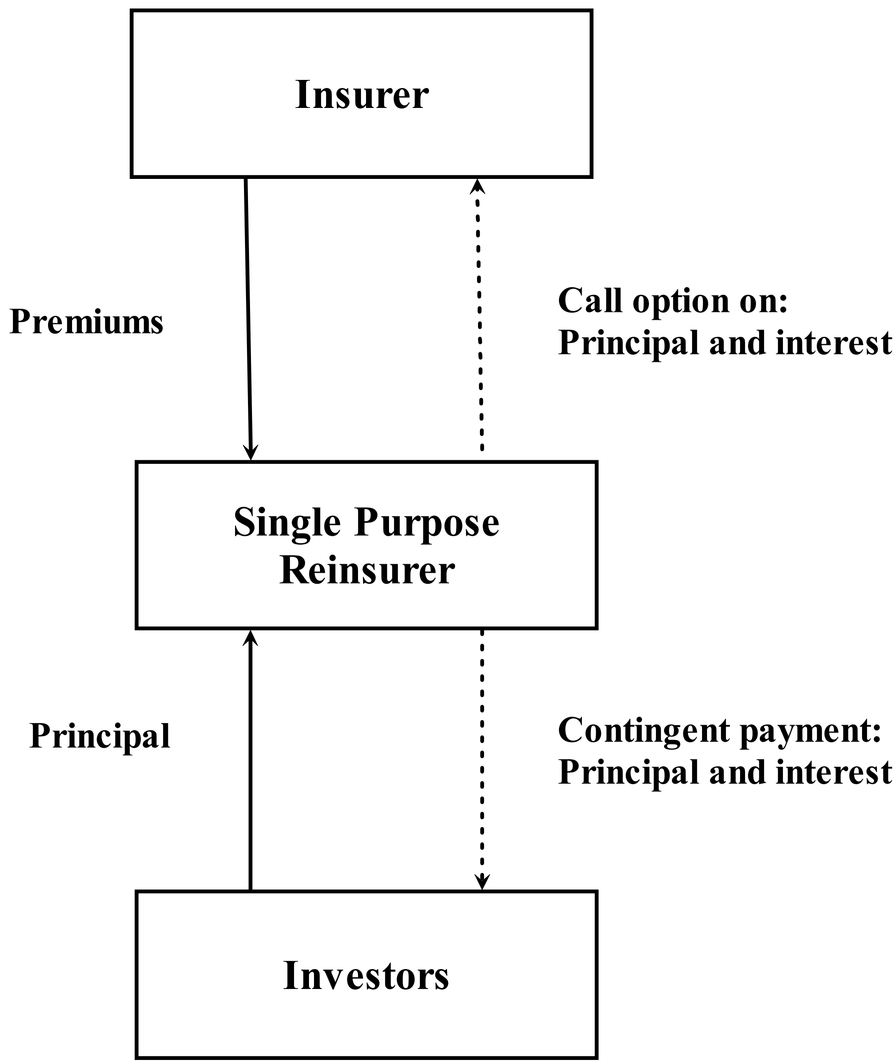

Figure 2

Simulating Insured Losses Using the AIR Model

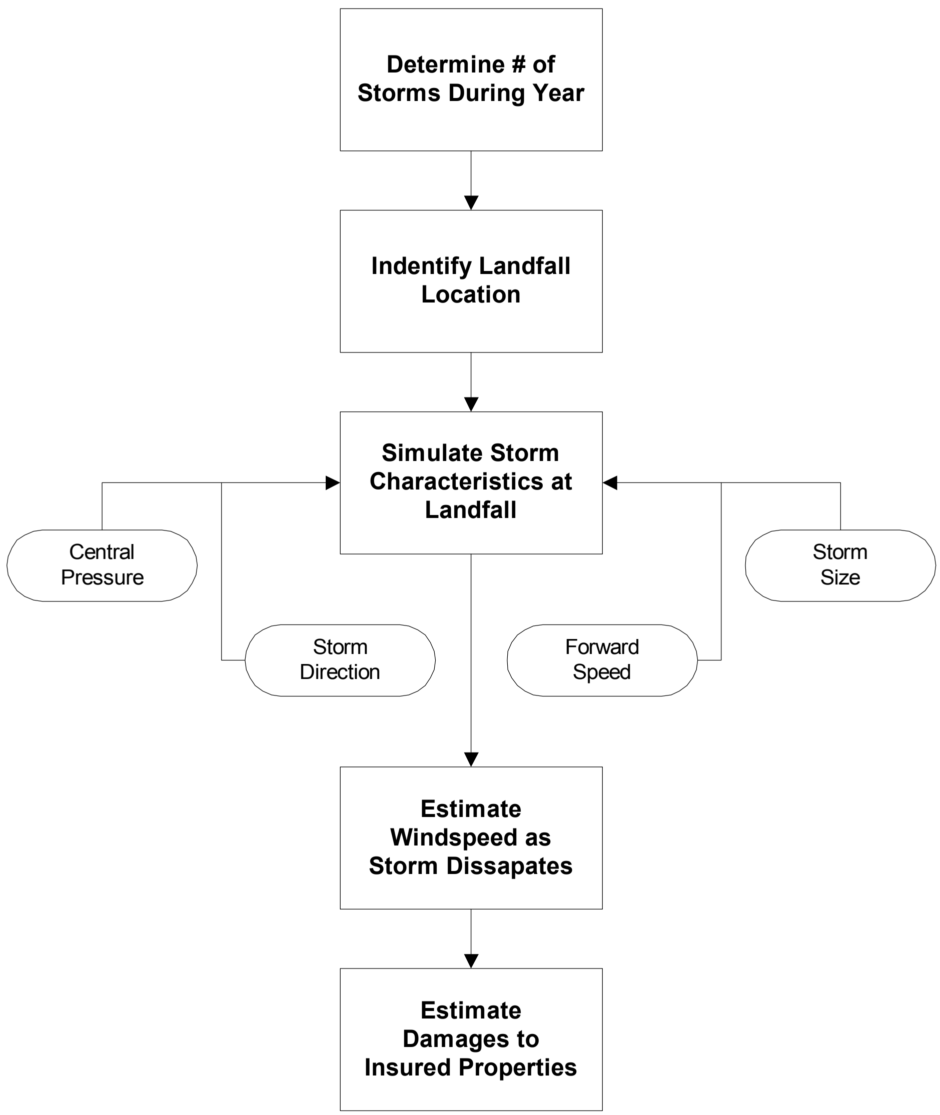

Figure 3

Company A: Large National Personal Lines Insurer

Panel A: Distribution of Company Exposures

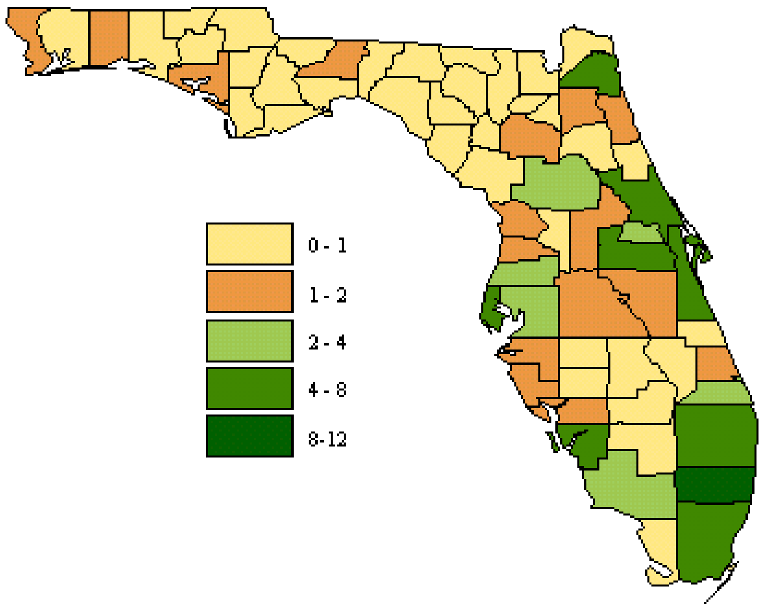

Panel B: County Market Share

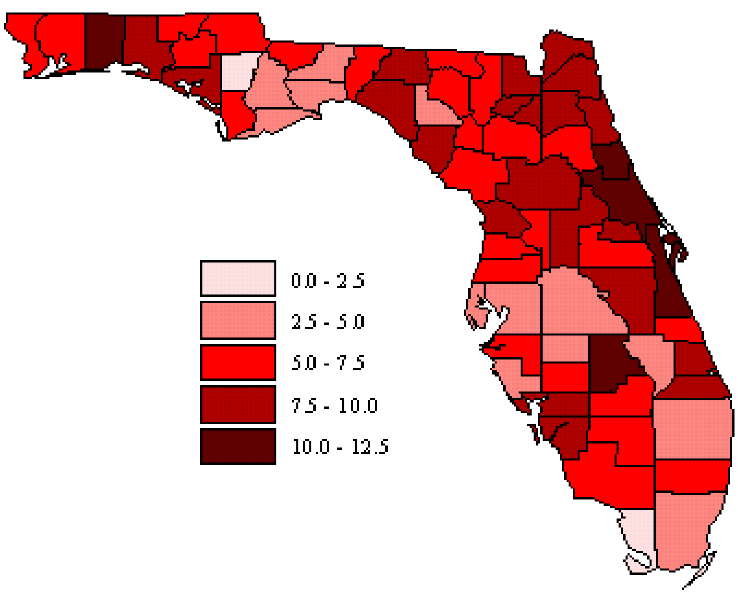

Figure 4

Company B: Quartile 1 Regional Florida Insurer

Panel A: Distribution of Company Exposures

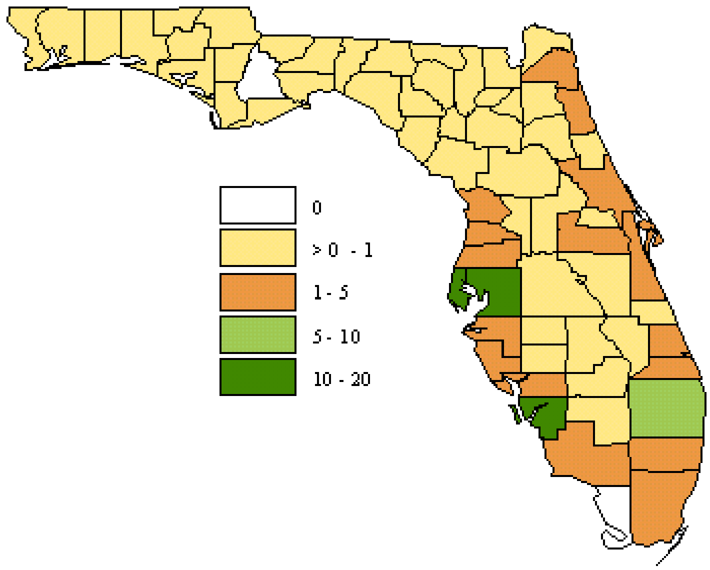

Panel B: County Market Share

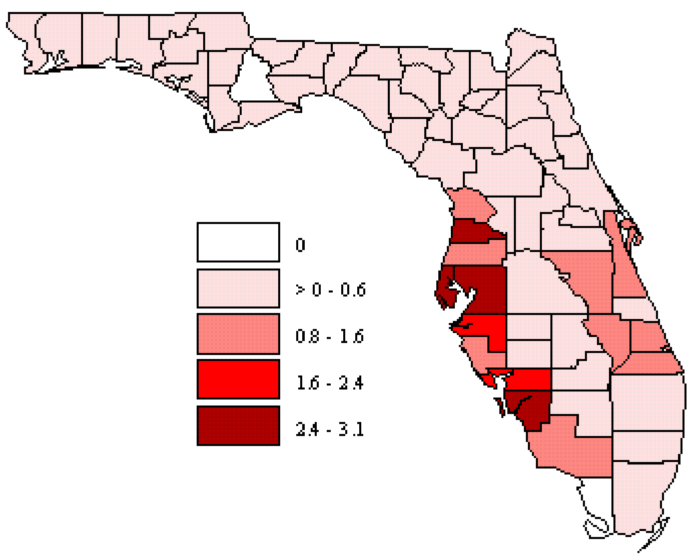

Figure 5

Company C: Quartile 2 Regional Florida Insurer

Panel A: Distribution of Company Exposures

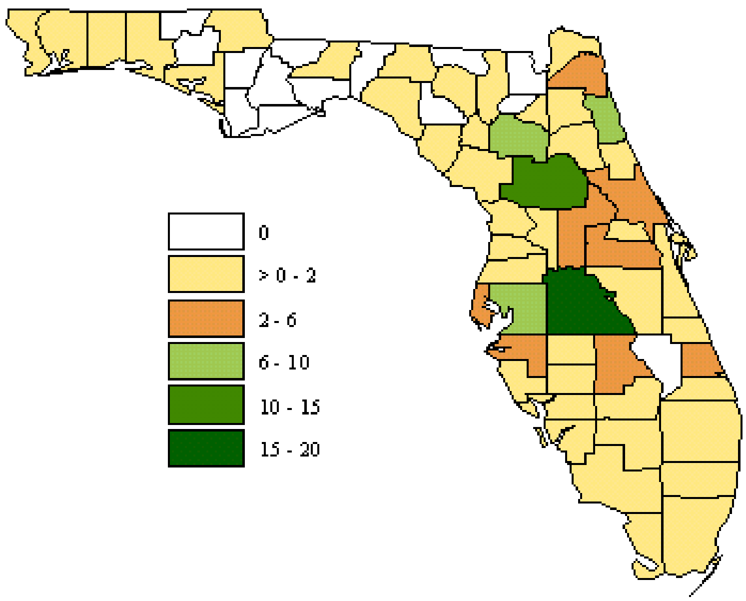

Panel B: County Market Share

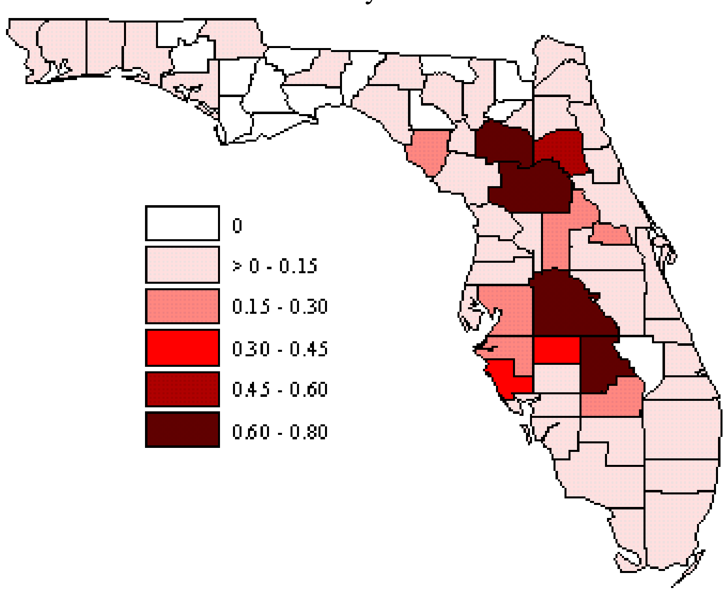

Figure 6
Variance Reduction Frontiers: Quartile 1 National Personal Lines Insurer (Company A)

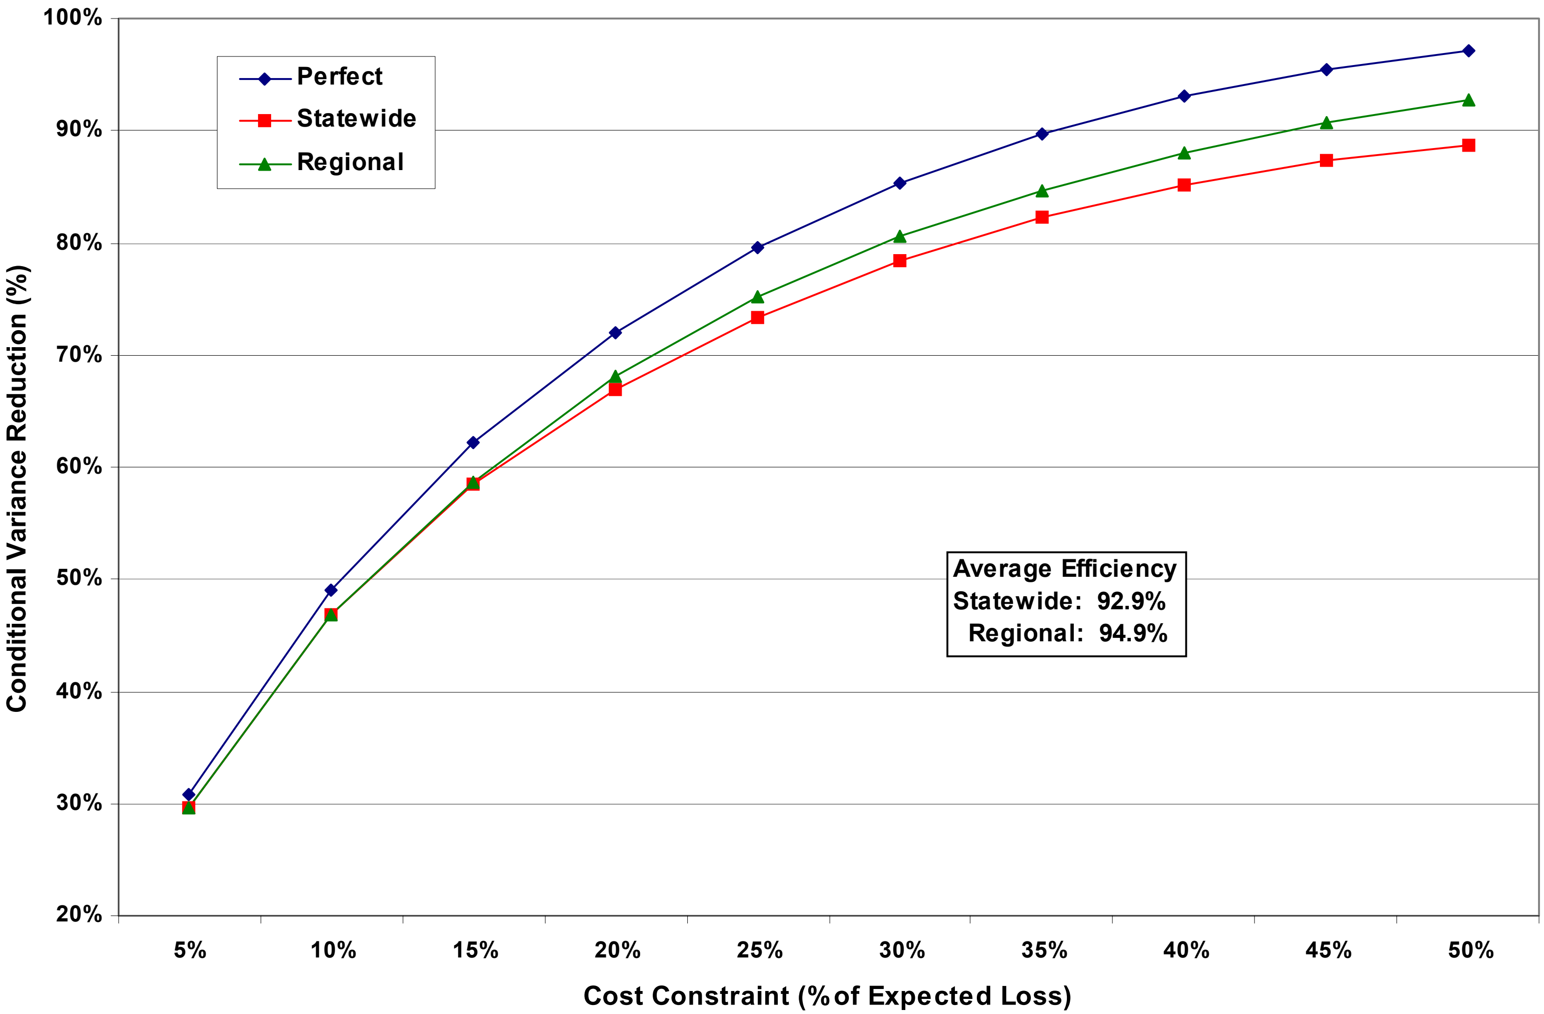

Figure 7
Variance Reduction Frontiers: Quartile 1 Florida Regional Insurer (Company B)

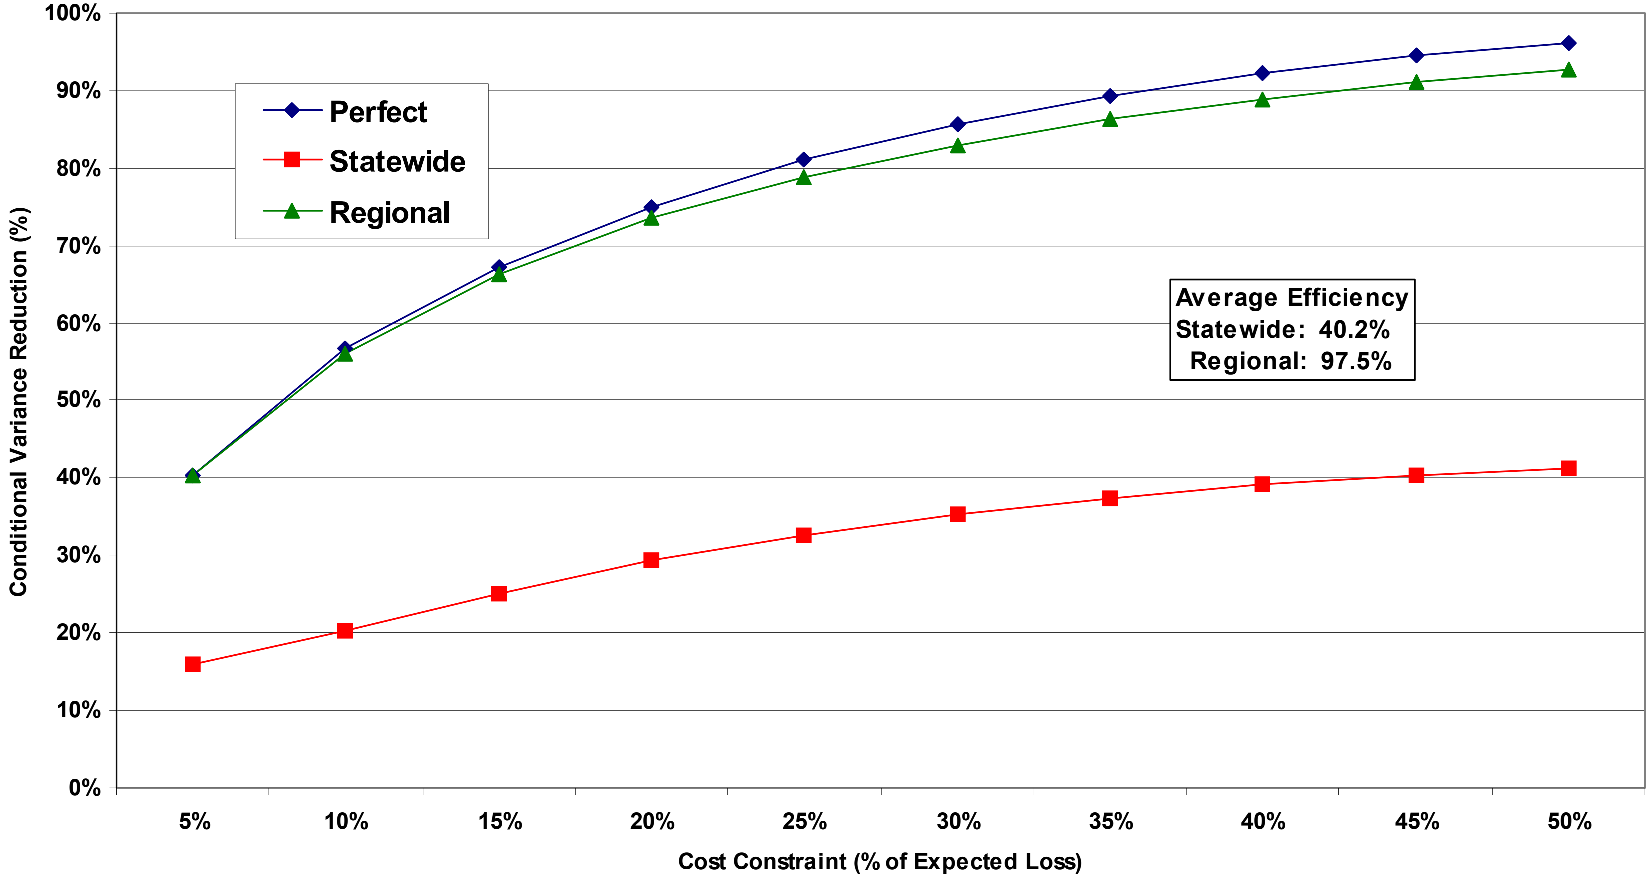

Figure 8
Variance Reduction Frontiers: Quartile 2 Florida Regional Insurer (Company C)

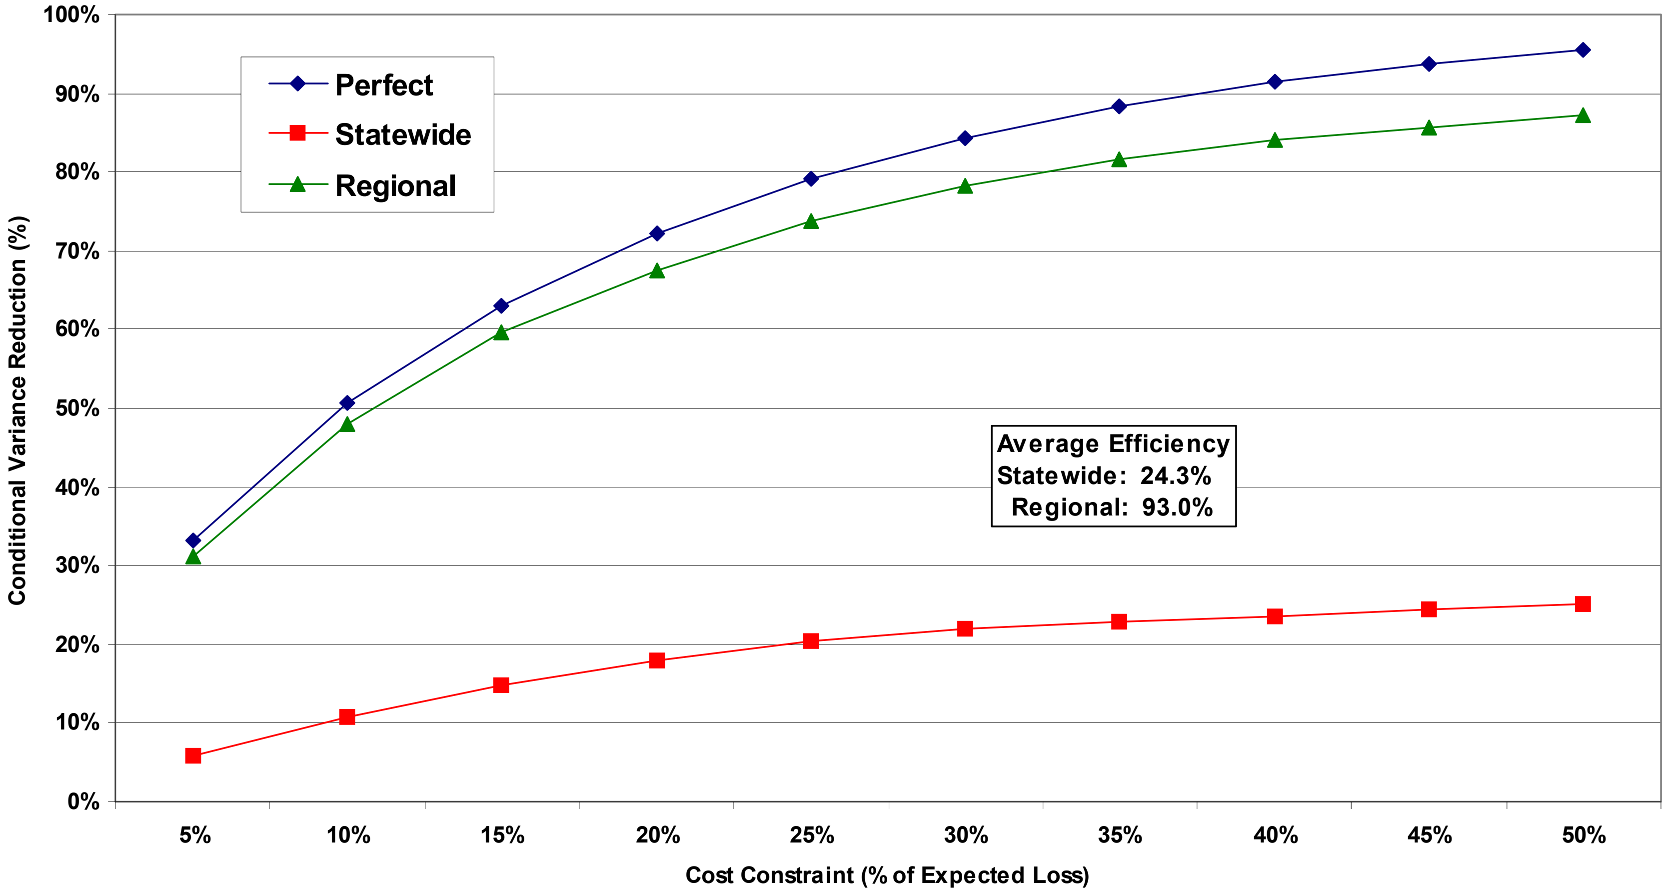

References

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

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

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

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

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

Cummins, J. David, David Lalonde, and Richard D. Phillips, 2002, "The Basis Risk of Index-Linked CAT Loss Securities," working paper, Wharton Financial Institutions Center, Philadelphia.

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

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

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

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

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

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

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

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

Addison-Wesley.

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

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

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

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

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

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

Swiss Re, 2001, "Capital Market Innovation in the Insurance Industry," *Sigma*, no. 3 of 2001 (Zurich, Switzerland.

Swiss Re, 2002, "Natural Catastrophes and Man-Made Disasters in 2001: Man-Made Losses Take on a New Dimension," *Sigma*, no. 1 of 2002 (Zurich, Switzerland).