Optimizing a Multi-Season Catastrophe Reinsurance Program With Private and Public Components

Rade T. Musulin, ACAS, MAAA and
John W. Rollins, FCAS, MAAA

OPTIMIZING A MULTI-SEASON CATASTROPHE REINSURANCE PROGRAM WITH PRIVATE AND PUBLIC COMPONENTS

By
Rade T. Musulin, ACAS, MAAA
John W. Rollins, FCAS, MAAA

Submitted to the 2001 Casualty Actuarial Society Call for Reinsurance Papers

Abstract

As risk-bearing organizations and the public at large have recognized the long-term frequency and potential magnitude of natural disasters, risk management solutions have emerged from both private enterprise and the public sector. A primary property insurer should choose a combination of ceded reinsurance, internal capital, and capital market products which minimizes its expected total cost of insuring consumers against catastrophes over a multi-year horizon, subject to risk tolerance and operating constraints. For this problem, the traditional metric of gross single event "probable maximum loss" is inadequate; a standard considering "probable aggregate retained loss" over the decision horizon is better.

The development of Dynamic Financial Analysis (DFA) to date has not given sufficient attention to natural disasters, the greatest threat to the solvency of many carriers. This paper shows how a property insurer can use commercial catastrophe models, public databases, some DFA concepts, and desktop computing tools to overcome two obstacles to integrating internal, private market, and public sector capital sources into an optimal program:

- Uncertainty about the market-wide impact of event frequency and severity over several seasons, which determines the response of the public reinsurer to the individual insurer.

- The complexity of the interactions between the responses of private and public programs to event experience.

The concepts and methods presented in this example may be generalized to solve a host of reinsurance program design problems involving multiple disjoint risk transfer vehicles.

The implications for the optimal strategy of the insurer of the decisions of public policy makers and the insurer's level of commitment to the market will be a backdrop to the analysis. Integration of the results into enterprise-wide DFA will be considered.

[1] Introduction: The Evolving Focus of Catastrophe Management

Though catastrophic events have occurred since the dawn of time, a prospective approach to managing their consequences is relatively new to mankind. The spectrum of pre-event actions that may be taken may be broadly subdivided into two groups: risk mitigation and risk distribution. Historically, risk mitigation came first – for example, the Romans built aqueducts and levees to manage floods. Risk distribution evolved later, with the development of probability theory and the ensuing rise of post-Renaissance fire and marine insurance clubs [Bernstein 1996]. We begin by stressing that both devices are still crucial to comprehensive public policy for catastrophe management. In this paper, we will focus solely on the optimal utilization of several political, scientific, and actuarial objects that have recently emerged in the risk distribution milieu, from the point of view of a central actor in it – the primary property insurance company.

Our protagonist is an insurer whose policy portfolio consists largely of property risks, and who is domiciled and writes business only in Florida. Florida-only property insurers held a very low statewide market share as recently as 1991, but are the most influential carriers in the market today. The parameters of the natural disasters, the insured loss exposures, and the insurance mechanisms we discuss are unique to Florida, but the form of the solutions and underlying reasoning are highly transferable to risk distribution problems in unrelated markets. Nonetheless, it is necessary to set the stage with a short discussion of the evolution of attitudes toward catastrophe management and development of scientific and political tools to facilitate risk distribution in this market in the 1990's.

In Florida, many political, economic, and climatological trends built toward a critical mass in the post-World War II era. These included an unusually low frequency of tropical cyclones, rapid population growth in highly vulnerable areas, and construction practices unable to withstand major hurricanes. During this period, insurers and regulators implicitly sanctioned this alignment of risk factors by failing to advance the state of the art in catastrophe risk assessment and ratemaking [Musulin 1997]. Citizens, insurers, and government largely ignored the volatile combination until the time bomb was "detonated" by Hurricane Andrew in August 1992. Afterward, a total rethinking of the methods for funding severe hurricane losses commenced among all the parties - an effort facilitated by major advances in measurement techniques for meteorological phenomena, structural damageability, and insurance losses. Over the next several years, many devices emerged which are currently the cornerstones of the functioning state insurance market. Those most crucial to our discussion are:

- Several competing catastrophe simulation tools (cat models) that allow risk-bearers and regulators to measure losses from future events on the basis of current coverages, demographics, and construction techniques, rather than historical experience.

- A competitive private market for the reinsurance of catastrophe risks, dominated by non-U. S. firms that specialize in this business and support the risk largely with retained capital.

- A public reinsurance facility for catastrophe risks, the Florida Hurricane Catastrophe Fund (FHCF or Cat Fund), which is an arm of the state and supports the risk with a combination of tax-exempt cash accumulation and contingent (post-event) securitization backed by sovereign assessment authority.

- Two residual markets for property insurance, whose claims-paying capacity are highly dependent upon capital from the FHCF, lines of credit, and post-event securitization backed by assessment authority: the Florida Residential Property and Casualty Joint Underwriting Association (FRPCJUA) and the Florida Windstorm Underwriting Association (FWUA)[1].

- Scores of Florida-only property insurance companies, some of which are affiliated with and backed by large national primary insurers, and others that were start-up companies backed by outside private capital but domiciled in the state.

We will review the workings of private catastrophe reinsurance and introduce the public reinsurance mechanism shortly. The first three items in the list above are important because a goal of the modern Florida property insurer, operating as a going concern, is to use the output of cat models to design an integrated catastrophe management program. The program must use public reinsurance, private reinsurance, and retained components as necessary to provide comprehensive risk distribution at minimum long-run cost, subject to an assumed desire to renew policies over the long run – even after a major event or "bad season".

Single-state companies are currently the predominant providers of property insurance in Florida. This fact is reflected in the last two statements in the list above, and does not square with academic expectations. Prevailing actuarial theory suggests that diversification of exposure among policies with weakly or even negatively correlated catastrophic loss experience distributions maximizes the profitability of the insurer, subject to constraints on the volatility of its aggregate experience, over any significant time horizon [Meyers 1996].[2] In plain English, each insurer should spread its policies geographically to the extent possible. Therefore, the Florida-only insurer should become uncompetitive over time due to substituting a more costly method of satisfying its net volatility constraint (e.g. reinsurance priced at market on concentrated business). The residual markets are state-specific creations, but still face the same exposure management problem as a private insurer (only with more constraints on action).

[1] The FRPCJUA was created shortly after Hurricane Andrew, but the FWUA was actually formed in 1970 for the purpose of covering only the peril of wind in the Florida Keys. The FWUA still covers only the wind peril, but has expanded significantly since 1992.

[2] Though Meyers framed his problem in terms of profit maximization, and we choose to discuss a cost-minimization problem, economists have shown that under very general conditions the two optimization problems have the same solution. In this paper, it is more to the point to present a cost-minimization problem which does not depend on the demand for primary insurance.

The public policy decisions which led to the bulk of Florida property exposure being carried by undiversified companies, whether in the voluntary or residual markets, reflect some interesting economics and have significant ramifications for all citizens of the state. We argue this point more fully in Appendix A. The implication of our argument, and a thesis of the paper, is that it is ultimately in the public interest for each single-state insurer to optimize its catastrophe funding strategy by using private and public capital to minimize its own long-run costs. As reinsurance is the most widely used vehicle for distributing catastrophic loss, it is also in the public interest to present ideas for determining the optimal reinsurance structure of the primary company. The rest of this paper presents an experimental case study on the challenges involved and some practical solutions. Along the way, we unearth some key insights about:

- The economics of catastrophe risk funding;
- The role of public sector capital;
- The measurement obstacles to complex funding solutions;
- The multifaceted nature of the models needed to analyze catastrophic cost distributions;
- The synergy between our model and traditional enterprise risk analysis tools;
- The public policy implications for stakeholders inside and outside the insurance industry.

These are summarized in the final section of the paper.

[2] Statement of the Problem - I

Fundamentally, our problem is one of economic efficiency. Stated qualitatively, the primary insurer wishes to minimize its *expected total cost* of insuring consumers against *catastrophic event losses* over a given time horizon, subject to a *risk tolerance constraint* over that time horizon.

The italicized terms above require definition in order to pursue a quantitative statement of the problem. A *catastrophic event* is a phenomenon bounded in time and space which contains an aggregate direct loss potential to the insurer exceeding its ability to fund the event through allowable premiums over the given time horizon. It almost tautologically follows that historical data is of insufficient credibility to derive event (and aggregate season) loss distributions. This condition warrants simulation of events in a catastrophe model, producing event loss data which can then be aggregated into event risk distributions. When the physical phenomenon is applied against the exposure of the insurer by a catastrophe model, an *event loss* is the dollar amount determined. The starting point for reporting event losses is usually "ground-up" insured damage, from which other functions such as "gross loss", "subject loss", and "net retained loss" can be determined with knowledge of the insurer's risk distribution program.

Risk tolerance constraints are a formal way of saying that the true definition of "catastrophic" is relative and depends on many parameters outside those describing the event itself. To evaluate catastrophe exposure against the insurer's appetite for volatility, several steps are required. Event losses must be translated into net losses for the insurer, with associated probabilities. Net losses must be accumulated over a given time horizon, the probability distribution of the aggregate retained losses must be computed, and the total cost distribution must be tested against the stated tolerance for risk.

More formally, we define a *risk tolerance constraint* to be the following type of statement: "Obtain less than a 100·P% chance of losing more than 100·K% of our policyholders' surplus over a period of N years", or:

$$\mathbf{Pr[C(N) > KS] < P} \qquad (1)$$

where P, K, and N are given by management, C(N) is the total cost function, and S is the insurer's wealth at time zero. The probability generated by given values of C, N, K, and S is the *risk tolerance function.*

The *total cost* function is stochastic in insurance. For the moment, let us reflect the risk tolerance functions of the actors as costs of capital in the total cost function. With no external risk distribution, the expected total cost of insuring consumers is simply the expected aggregate direct losses over the N year horizon, plus the opportunity cost of internal risk capital (whether provided by stockholders, policyholders or others) committed to the catastrophic exposure; the collective risk representation is:

$$E[C(N)] = E[X_1 + X_2 + \ldots + X_{M(N)}] + \theta_0 \qquad (2)$$

where both M (the number of events in N years) and each X_i (individual event loss) are random variables, and θ_0 is the opportunity cost of internal risk capital. With a reinsurance program, the total cost becomes the sum of the premiums paid to reinsurers over N years, plus the expected aggregate retained losses and cost of internal capital committed to retained losses (θ^*), or:

$$E[C(N)] = \pi + E[R_1 + R_2 + \ldots + R_{M(N)}] + \theta^* \qquad (3)$$

where π is a guaranteed-cost reinsurance premium[3] and R_i (the retained event loss) is a potentially intricate function of X_i which depends upon the structure of the reinsurance program. The external reinsurance premium in turn reflects expected ceded losses, a cost of external capital committed to the ceded losses (θ'), and transaction costs (τ):

$$\pi = E[\{X\}] - E[R(\{X\})] + \theta' + \tau \qquad (4)$$

[3] We ask readers familiar with catastrophe reinsurance to let us ignore reinstatement premiums for the moment. Reinstatements are incorporated into the problem later, as an increased co-participation on the ceded losses (thereby raising $R(X_i)$). They could instead be formulated as a conditional addition to reinsurance premium (thereby making π stochastic). In most treaties, estimated premiums are adjusted over the term in conjunction with changes in estimated exposure, but again we simplify things for presentation.

In an equilibrium distribution of risk, we expect $(\theta' + \tau + \theta^{*})$ to be at a minimum and below θ_0. In other words, we presume that the primary insurer will buy reinsurance to the point at which the marginal external cost of capital plus transaction costs embodied in the reinsurance premium exceed the marginal internal cost of capital for self-funding, or (substituting (4) into (3) and taking the total static differential):

$$|\Delta(\theta' + \tau)| = |\Delta\theta^{*}| \qquad (5)$$

The rates of change are not equal except in equilibrium, and likely depend on the correlation of the catastrophic exposure with the other risks in both the primary insurer's and the reinsurers' portfolios, as shown by Kreps [1990] and others. Now consider the possibility of public reinsurance, so that the reinsurance premiums paid to private (A) and public (B) sources comprise differing costs of capital and transaction costs:

$$\pi = \pi_A + \pi_B \qquad (6)$$

We explain in §[3] below why it is possible to set π_B with a cost of capital component below private market rates and internal rates, so that real savings may be achieved and internal capital freed up by use of the public facility. The underlying rationale for state-backed reinsurance is that the sovereign authority of the state can be used to defer part of the capital commitment until after an event, thus lowering the required cost of capital component of π_B. Recalling (5), then, the introduction of public capital implies an opportunity to establish a new equilibrium reinsurance structure with lower total costs.

Rather than model the cost of capital in the total cost function, we consider it implicitly through development of the risk tolerance constraint. Such a formulation is much easier to work with experimentally, when event gross loss data is provided but companywide hurdle rates and other cost of equity measures are unavailable. The company optimizing reinsurance strategy in the context of an enterprise-wide DFA model may wish to proceed with minimization of the total cost function, including the costs of capital determined elsewhere in the model but dropping the explicit risk tolerance constraint. We suppress costs of capital in our working equations for the experiment while acknowledging the importance of the concepts.

When dealing with a defined reinsurance structure later, we will find it convenient to separate the expected aggregate retained losses (EARL) into components based on the architecture of the program (retention, R_r; co-participation, R_c; and excess, R_x) rather than by individual event:

$$\text{EARL} = E[R_r(X_1 + \ldots + X_M) + R_c(X_1 + \ldots + X_M) + R_x(X_1 + \ldots + X_M)] \qquad (7)$$

We will also separate reinsurance premiums paid to private and public sources, so our logic in the experimental simulation exercise will be based on:

$$E[C] = \pi_A + \pi_B + E[R_r + R_c + R_x] \qquad (8)$$

In the experiment, R_c will be further subdivided into the participating losses in every layer.

The reader must be careful not to conclude that replacing private with public reinsurance priced at a lower cost of capital is risk-free. This is not true; in fact, a thesis of the paper is that use of the public facility generates a measurable change in the insurer's risk tolerance function. Therefore, our focus is not on arbitrage (improving costs without taking risk), but on managing the trade-off between reinsurance premium savings and increases in the probability of excessive retained losses over the decision horizon. The introduction of public capital simply mitigates the cost effect of changing expected retained losses (and therefore the risk tolerance function).

In our research, P, K, N, and S are given, and {X} is provided by a sample data set from a catastrophe model. Calculation of the other variables (reinsurance premiums and retained losses) requires specification of a company reinsurance structure. We turn to some background on both public and private reinsurance markets in order to develop reasonable experimental assumptions for discussion.

[3] The Public Reinsurer – the Florida Hurricane Catastrophe Fund

In the wake of Andrew it became obvious that the capital required to fund potential losses on the Florida housing stock accumulated during the construction boom of the 1946-1991 period was far beyond that which was available from sources extant before the event. An immediate and enormous infusion of capital was required in order to renew existing insurance contracts and provide capacity for additional population growth. Florida's public policy planners faced an unpalatable menu of options for accomplishing this [Musulin 1999]. The gap between public expectations of "affordable" coverage based on pre-Andrew premium levels and the reality of expected cost estimates based on information learned from Andrew was wide, and had to be bridged fast. After much debate, the state settled on a solution that combined price increases, coverage restrictions, and the contingent use of its sovereign authority to secure bond issues that would raise billions of dollars for claim payments, while deferring the total cost for decades after a major event.

The state had two advantages over private markets as a source of needed capital:

1. Under accounting conventions, insurers cannot issue post-event bonds without establishing a liability for their repayment, in effect forcing losses to be pre-funded. The state faces no such restriction, allowing it to create billions of dollars in capacity without the need to charge an "up front" premium.

2. A properly structured state program qualifies for exemption from federal income taxation on accumulated premiums. Premiums collected by private insurers over many years to cover infrequent catastrophic events face severely unfavorable tax treatment – the premium "income" from each storm-free year is fully taxed at

year's end, since no calendar year losses are available to offset it. In the catastrophic year, tax carrybacks and carryforwards generated by the devastating loss are of only limited effect in recovering the excess taxes paid in the past and likely in the future.

Following Andrew, the state enacted laws allowing three entities, the FHCF, FWUA, and FRPCJUA, to issue bonds to cover losses in excess of cash raised through premiums. While all three entities can affect a private insurer's net experience from a catastrophic event, we will confine our discussion to the FHCF.

Table 1 summarizes the pertinent characteristics of the FHCF. It has been able to provide a reinsurance product to insurers at a cost below that in the private reinsurance market because it does not charge a risk load[4] and because it can accumulate premiums free of federal income taxes. It can also issue tax-exempt bonds.

The attributes of the FHCF carry some unsettling implications for an insurer operating in Florida and wishing to responsibly utilize a public reinsurance program:

- The amount of coverage can vary significantly from year to year. It will drop in the year following an event, leaving a gap that will have to be filled by internal capital or reinsurance markets already stressed by the event.

- The change in capacity following an event will depend on <u>industrywide</u> losses, which may not be correlated with those of the primary insurer. This disconnect is sometimes called "basis risk".

- FHCF premium is constant[5], but its capacity varies considerably. Thus, an insurer may face widely varying rates on line (ROLs) from season to season for coverage even though the premium dollars are predictable.

- Florida rating law prohibits an insurer from charging its customers for reinsurance that "duplicates" coverage available from the FHCF, effectively forcing insurers to replace private coverage with public coverage as its capacity grows. This may make it difficult to pay the cost of replacing FHCF coverage after an event due to the lag in recovering additional reinsurance costs through the primary insurer's filed rate structure.

---

[4] Since much of the capital funding an event is only committed after it, the opportunity cost of capital is lower, as explained in the previous section.

[5] In this context, the term "constant" does not mean constant absolute rates. The FHCF reviews its rates annually and makes some minor changes, but these are related only to desired, not actual, amounts of coverage provided.

Table 1 – Key Characteristics of the FHCF [F.S. §215.555]
• It is a reinsurer covering residential policies written by primary insurers operating in Florida.
• Participation by primary insurers is mandatory.
• It covers buildings, appurtenant structures, and contents, but not additional living expense (Coverage D in most property policies).
• It enters into contracts of adhesion (meaning that insurers cannot negotiate terms) with all insurers writing residential coverage in Florida. The contracts specify an insurer's event retention and season aggregate coverage amount.
• Coverage is direct, without regard to an insurer's private reinsurance arrangements, and limited such that an insurer cannot collect more than 100% of a loss from the FHCF and private sources.
• It charges premiums based upon the insurer's exposure (not subject premium!) at a rate per $1,000 of coverage. Rates vary by ZIP code, construction, deductible, and several other attributes.
• Each insurer may choose a coverage level of 45%, 75%, or 90%.
• Its premiums are based on desired industry aggregate coverage of $11 billion per season, regardless of the actual coverage it provides. It sets rates based on long term expected loss costs discounted for investment income; its capacity varies based upon accumulated funds and bonding ability.
• The industrywide aggregate retention (the sum of the individual insurer retentions) per event is slightly over $3 billion, and indexed annually to the change in aggregate industry exposure.
• It funds losses through a combination of premiums paid pre-event by participating insurers and post-event bonds secured by ensuing assessments on all property-casualty premiums in the state excluding Workers Compensation. The sum of all outstanding assessments is limited to 6% (of written premium), and 4% is the limit on assessment in any one season.
• Its capacity in a given season is limited to the lesser of $11 billion, or its accumulated premiums plus the proceeds from bonds supported by the assessment authority remaining in the season of loss, as noted above. In 2001, its capacity is $11 billion.
• Unless the first season of loss is fully funded by accumulated premiums, in the year following a major loss its capacity will drop significantly due to the depletion of accumulated premiums and the limitation on total assessments. If the Cat Fund incurred a full $11 billion loss in 2001, its capacity in 2002 would drop to about $7 billion.
• Each insurer's share of the retention and capacity is based on its share of the total FHCF premium. This is accomplished through retention and coverage multiples, which set the insurer's retention and limit as a multiple of its FHCF premium.

- The FHCF does not cover non-residential commercial risks or additional living expenses on covered policies. This means the coverage provided by the public facility does not interlock well with private reinsurance, even if the private program is designed specifically to allow the FHCF to reinsure certain layers.

[4] Review of Private Reinsurance Markets

We assume the reader has a working knowledge of catastrophe reinsurance concepts and terminology, so our introduction to private markets will be succinct and focused on characteristics affecting our problem.

A variety of reinsurance programs and capital market solutions exist for ceding a primary insurer's catastrophe risk, and product innovation in this area continues at a rapid pace. Our purpose is not to survey, so we will focus on traditional layered excess-of-loss catastrophe reinsurance with cedant's retention calculated as a percentage of subject matter premium (SMP). Our experimental program will be designed to cover 95% of all losses in the reinsured layers from all property lines of business (excluding automobile).

Several attributes of private reinsurance programs are germane to our discussion:

- *Multiple "layers" of coverage*, each parceled among subscribing reinsurers in the layer, with the ceded premium rate per coverage unit or "rate-on-line" decreasing for higher layers.

- *Co-participation*, typically 5% of the losses in each layer, retained by the cedant. Layers are structured "thinly" enough, and co-participations small enough, that a moderate change in co-participation should not affect the rate-on-line charged for the layer.

- *Per-event coverage*, triggered by a qualifying event, with layering expressed in terms of single event loss, even though contracts are written for a season.

- *Reinstatement premiums*, which fund coverage for multiple events in a single contract season. After an event loss, coverage for a second (or third, or Mth) event is restored for an additional premium, usually computed as the original premium pro-rata to the amount of the limit used in the previous event. When constructed in this manner, reinstatement premiums can be restated as increased participations in program analysis. Generally, the optimal number of reinstatements allowed should decrease in higher layers less subject to frequent losses.

- *Potentially volatile rates-on-line* from year to year, reflecting individual loss experience and general market conditions. This is true even though market prices often reflect expectations of multi-season relationships between cedants and reinsurers.

- *Coverage of all losses* in primary property insurance contracts subject to the program (the reinsurer "follows the fortunes" of its cedant).

- *Coverage net of inuring* reinsurance, such as pro-rata property treaties (quota share and surplus share), facultative contracts, and public reinsurance.

Even forgetting the role of the FHCF for a moment, private catastrophe programs pose significant financial analysis challenges. Most financial analysis of catastrophic events focuses on the impact of a single-event "probable maximum loss" (PML) on the insurer's operations. This approach is inadequate because it:

- Ignores the problem of an accumulation of retained losses due to moderate events, which can produce a greater probability of significant surplus loss over a multi-year horizon than the impact of "the big one".

- Ignores the possibility of multiple events in a season exhausting the reinsurance coverage in some layers.

- Fails to adequately consider the post-event changes in the cost of internal capital and/or external reinsurance.

- Fails to consider the insurer's ability to renew policies as a going concern, either due to changes in underwriting capacity or the inability to pass higher capital costs on to policyholders because of rate regulation.

A program including the FHCF faces even more thorny problems because the public coverage structure differs from that of traditional reinsurance. The nature of other current and proposed government programs, perhaps backed by Federal reinsurance[6], is likely to produce similar complications in designing an integrated public/private ceded reinsurance program. Considering the FHCF, for example:

- There is only one layer, with one ROL, because coverage is stated in terms of a season aggregate amount with no limit on the number of events covered. This is true even though coverage is triggered by, and retentions are charged by, event.

- Only some property losses are covered - the FHCF does not cover non-residential lines.

- "Additional living expense", a.k.a. "coverage D" loss, is not covered. This can create particular problems if an insurer has a policy level deductible applicable to losses from all coverages.

- Coverage is direct, and limited such that an insurer cannot collect more than 100% of an event loss from private and public sources. This forces the primary insurer to carefully coordinate coverage to prevent an inadvertent limitation of FHCF payout. For example, an insurer writing large risks and making extensive use of facultative reinsurance could run afoul of this limitation.

- Premiums are constant, but coverage may drop considerably after an event, and the degree of drop-off varies from year to year according to the accumulated experience of the FHCF since its inception.

- Coverage over multiple seasons is dependent on several external variables. Some parameters are political, such as the assessment limitations for debt service, and some are given by the capital markets, such as the interest rates at which the FHCF can issue tax-exempt bonds on the open market following an event.

[6] Federal natural disaster legislation has been advanced in recent years through vehicles such as H.R. 21, though not yet passed into law. Many of the proposals would auction federal reinsurance contracts to both private and public insurers, and some would give preference to state-backed catastrophe reinsurers such as the FHCF.

In the following sections of the paper, we shall demonstrate how catastrophe models and reinsurance program simulation can be used to overcome these problems and design a catastrophe risk distribution system which meets the objectives and constraints outlined in §[2].

[5] Statement of the Problem – II

We will bring specificity to our problem by constructing an experiment. Exhibit 1 introduces the current private reinsurance program for A-Florida Insurance Company, a hypothetical single-state property insurer. Several parameters fixed by management or the market are assumed.

A-Florida has an opportunity to introduce public coverage into its reinsurance program next year; relevant parameters of the public facility are shown at the bottom of Exhibit 1. The actuary notes that the public reinsurer can provide coverage excess of $15 million per event, with a season aggregate payout of up to $50 million. For this coverage, the rate-on-line of 4% is cheaper than the average rate-on-line in the private program across the layers where the public reinsurer can provide coverage. Therefore, A-Florida believes it can reduce its expected total cost of catastrophe risk distribution over a multi-year horizon, while making an acceptable change in its risk tolerance function, by replacing some of its private cover with the FHCF.

Operationally, A-Florida does not wish to totally cancel coverage in layers likely to be covered by the Cat Fund, because to do so would forego a market relationship, with benchmark pricing, in the event of an abrupt drop in public capacity.[7] For this reason, it devises a program as follows:

1. Continue buying some private coverage in layers that overlap with the new FHCF coverage (e.g. excess of A-Florida's FHCF retention by event and within the bounds of the season aggregate coverage amount), but cut back coverage levels (meaning the complement of co-participation rates) in these layers.

2. Buy public coverage and retain the first right to public program recoveries from an event to fund three costs: event losses in excess of the highest layer of private coverage (over-the-top losses R_x), reinstatement premiums (treated as reductions in private recoveries) in each private layer, and "gap" losses, defined as the difference between the original coverage level (95%) and the revised (cut-back) coverage level in each applicable private layer. Gap losses are part of R_c.

[7] Protection against public coverage failures may be obtained after the fact on the spot market for private reinsurance, but at market rates inflated by the surge in demand for capacity after a stormy season. It is outside the scope of this paper to construct such an experiment and model the expected increase in π_A in the years after an event, though we suspect the solution of conditional reliance on the spot market to be severely sub-optimal with respect to C(N) where N>1. Our tools could be adapted to solve such a problem.

3. Provide private reinsurers with a contractual right to the remainder of the Cat Fund recoveries, shared pro-rata to the ceded losses for each reinsurer, in exchange for a premium credit. A-Florida expects that, after an event, it will be able to renew the existing coverage and restore it to 5% co-participation, with the premium credit reduced or eliminated to compensate for any drop in public coverage.

Therefore, the choice variables in the experiment are:

- The coverage level selection (complement of the co-participation rate) for the FHCF.
- The coverage level selections in each reinsured layer of the private program.

We select this type of program alteration, rather than elimination of private layers or horizontal changes in layering or retentions, based on the business reasoning above and the presumption of equilibrium explained in §[2]. The total amount of coverage provided is assumed optimal, and there is no total cost savings to be generated solely from changing layer boundaries or retentions. If the layer widths are not changed, the rate-on-line should be constant in each layer almost regardless of the co-participation selected, allowing easy (linear) measurement of the premium savings from the substitution of public for private coverage. Under these assumptions, the most reasonable choice variables are private and public coverage levels.[8]

If the public coverage was guaranteed at a below-market rate for N years, and it covered exactly the same subject losses as the private program, the problem would now be simple – buy the maximum public coverage level and use it to replace private coverage dollar-for-dollar, starting with the lowest (most expensive) applicable layers. Unfortunately, neither premise is true.

Recall that the scope of public coverage is restricted (commercial non-residential property and Coverage D losses are not subject), and that the availability of public coverage after year 1 is at risk due to reductions in its cash balance and borrowing capacity. Cat Fund performance in years 2-N depends upon its cumulative experience in prior years of the horizon. We must devise a way to analyze, over the entire planning horizon, the tradeoff between the reduction in reinsurance premium ($\pi_A + \pi_B$) from including public coverage and the increase in expected aggregate retained losses $E[R_r + R_c + R_x]$ from public coverage inequities as well as potential coverage failures in later seasons.

Several custom tools are required to conduct the investigation:

[8] In a purely private program, the tools introduced in our experiment could be used to determine the optimal number of reinstatements allowed in each layer, reflecting the effect of loss frequency on the risk tolerance constraint. We believe this risk often receives scant consideration in private program design. Here, we assume that the insurer has already selected an optimal number of reinstatements by layer, and these provisions will not be choice variables.

- A financial model of the public reinsurer tracking both its cash position and bonding capacity, and calculating potential amounts of coverage failure in each season over N years of industrywide event experience.

- A database with records for every event simulated by the cat model, augmented with several fields beyond A-Florida's gross losses: company losses subject to private reinsurance, company losses subject to public reinsurance, and industrywide losses subject to public reinsurance.

- An integrated model of the proposed private and public reinsurance structure of the insurer which allocates gross event losses to each component of EARL, net of FHCF recoveries, as well as the experience of each reinsurer net of its share of FHCF recoveries.

- A computer simulation engine which randomly generates seasons of hurricane experience for both the insurer and the industry from the event-level catastrophe impact database, runs the public reinsurer financial model to determine FHCF coverage, feeds the experience into the proposed reinsurance structure, and accumulates event losses in each "cell" of the program (whether retained or ceded). A cell is defined as a constant (vertical) proportion of a given (horizontal) layer of coverage.

- A database and spreadsheet summarizing the results and calculating statistics to evaluate the expected total cost function and the probabilities associated with the risk tolerance constraint.

Figure 1 below is a heuristic diagram of the tools and their relationships in the experiment. We will explain our construction of each of these in turn.

Figure 1: Diagram of Workflow – A-Florida Experiment

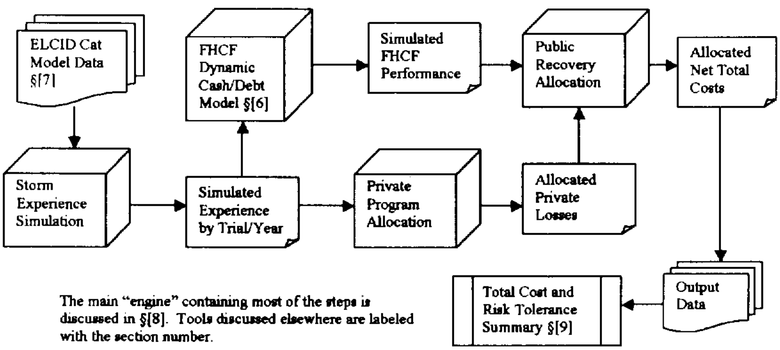

[6] Dynamic Analysis of the Public Reinsurer

When public reinsurance is incorporated into a layered private program as described above, determining retained losses R(X) from gross losses is non-trivial. We must design a model that uses industry-wide event losses subject to the FHCF (which in turn depends upon the gross losses to the primary insurance market in aggregate from each event) to determine the effect of catastrophes upon its cash balance as well as its debt burden. Catastrophe impacts on each of these accounts must be accumulated and carried forward to each year in the experiment horizon to determine the coverage available to the industry in later seasons.

Our dynamic model for the cash account of the FHCF is the following difference equation:

$$\mathbf{S_t = S_{t-1} + \pi_t + I_t - O_t - L_t} \qquad (9)$$

Cash at season's end (S_t) equals beginning balance (S_{t-1}), plus reinsurance premiums received from primary insurers (π_t), plus investment income on beginning balance (I_t), less operating expenses (O_t), less losses paid from the cash account (L_t) [9]. L is the random variable which depends on the hurricane experience (losses subject to Cat Fund) in the current year as well as the cash position prior to the hurricane(s): ($S_{t-1} + \pi_t + I_t - O_t$).

The coverage commitments of the FHCF far exceed its current cash balance. If the subject losses exhaust the cash balance, it must borrow money on the open market to honor its coverage. Its capacity to do so depends on its authority to assess Florida consumers (through their insurers as collection agent) to pay its creditors. Therefore, modeling the "assessment authority account" is the key to determining the probability of, and expected amount of, coverage failure in a bad year. Opening assessment authority for the year is determined by the model according to the rules in §[3] regarding single-season and aggregate caps on assessment amounts, and then reduced by the per-annum amounts of current assessments required to service all outstanding debt issues (tracked by the engine). The remainder is available to support new debt – once a term to maturity and a market interest rate are stated, simple financial calculations determine the bonding capacity. If the bonding capacity plus the cash account is short of the subject losses for the year, we have a coverage failure.

When a coverage failure occurs, all insurers participating in the Cat Fund are penalized equally – each insurer's recovery is pro-rated to reflect the aggregate shortfall. For example, if $11 billion in coverage was promised to the industry (and allocated to each insurer based on its ceded premium) at the outset of season 2, but only $8 billion in cash and debt was available at season's end due to exhaustion of assessment authority in

[9] Balance sheets, investment yields, and expense information for the Cat Fund are available in its Annual Report, issued each season by the Florida State Board of Administration.

season 1, an insurer whose promised coverage amount was $220 million would only receive 8/11 of it, or $160 million. Hence, the summary statistic from each year in the dynamic model is the "industry recovery ratio", which is 100% unless a coverage failure occurs[10]. This ratio is used to calculate the coverage reduction in A-Florida's public reinsurance.

A template for a dynamic model of the cash and debt accounts of the FHCF is shown in Exhibit 2.

We assume a fixed market interest rate and term to maturity in our model. An obvious enhancement when integrating our research into an enterprise-wide DFA system would be to tap into the interest rate generator used elsewhere in the system, in order to calculate a stochastic market yield curve for FHCF tax-exempt debt each year[11].

[7] The Event-Level Catastrophe Impact Database

An excerpt from the final event level catastrophe impact database (ELCID) is shown in Exhibit 3. The key fields entering the simulation engine are

- Season index
- Event index (within season, reset to 1 every year)
- Company losses subject to private reinsurance
- Company losses subject to FHCF
- Industrywide losses subject to FHCF

Depending on the cat model and generation thereof, several intermediate steps may be required to augment raw output for ELCID:

- Indexing of the events within seasons, so that they may be kept in order in the simulation.
- Inclusion of time-element losses (such as additional living expenses), which are subject to the private program but not the FHCF, in the gross losses.
- Estimation of ceded losses to reinsurance which inures to the benefit of the private cat program (i.e. surplus share).
- Estimation of the commercial property portion of the gross losses.

---

[10] This is very much like the complement of the "expected policyholder deficit ratio" used in solvency analysis [Butsic 1994], but related to coverage instead of net worth.

[11] In a real bonding exercise, the FHCF would issue multiple debt instruments in various tranches to minimize its open-market borrowing costs, so the selection of a single average term to maturity is an oversimplification.

- Calculation of losses subject to private reinsurance by subtracting surplus share ceded losses from gross losses including time element.

- Calculation of losses subject to public reinsurance by subtracting commercial and time-element losses from the gross losses.

- Estimation of industrywide losses subject to public reinsurance by applying market share tables by line of business (homeowners, mobile home, commercial property) and county (territory)[12].

Some complex database logic is used to carry out the steps outlined. We begin by obtaining a data set containing 50,000 simulated years of company event experience, with gross losses by event, by county, and by line of business from the cat model. We populate a table containing estimated proportion of time-element losses by event size and line of business from a review of historical hurricane experience by coverage at the company, and apply it to the raw data to incorporate time-element losses. We use line of business categories to separate commercial losses from residential losses covered by the FHCF. We remove simulated surplus share losses (provided by the cat model output), subject to the aggregate cap on the surplus share program, from the event losses and calculate losses subject to private reinsurance. We also calculate losses subject to public reinsurance by removing the commercial lines of business and the time-element portion of the gross losses, but not the surplus share ceded losses.

Industrywide exposure by county for four business categories: residential, mobile home, commercial-habitational (which is covered by the FHCF) and true commercial (which is not covered anymore), is obtained from publicly available FHCF data. Grouping the company's business into the four categories, we derive a "market share factor" as the ratio of the company exposure to the industrywide exposure in each group/county combination. Finally, we divide company losses subject to FHCF by the market share factor to generate estimated industrywide losses by county and line of business. This leveraging method assumes the modeled damageability (catastrophe loss cost per unit of exposure) for A-Florida's book of business is similar to that for the industry in general. This is an imperfect assumption but one that allows us to build the dynamic financial model of the FHCF without obtaining its entire cat model results.

The end result is a data set sorted by season index and event index within year, showing A-Florida's subject losses and the estimated industry subject losses in total for each event simulated by the physical cat model. Indexing allows randomization and application of a simulation algorithm as described below.[13]

---

[12] Some catastrophe models can provide industry-wide gross event losses directly, but using that data would not eliminate the problem of determining losses subject to public reinsurance from the gross amounts.

[13] Special thanks go to Steve E. Wallace, Vice President – Operations of Florida Farm Bureau Casualty Insurance Company, for the programming which built the ELCID database.

[8] Randomization of Data and Simulation of Experience

The distribution of each component cell of the total cost function for A-Florida (and incidentally, the cost functions for each of its reinsurance layers) is estimated empirically through simulation. We build a spreadsheet in Microsoft Excel 2000 to store the ELCID data, house our dynamic model of the public reinsurer, store all the parameters by layer of the private reinsurance program, simulate random catastrophe experience season by season and trial by trial, and most importantly, apply the simulated experience simultaneously to the public and private reinsurance programs to gauge the true response of an integrated program to catastrophic events. The randomization and simulation logic is developed in Visual Basic for Applications using only Excel-based objects.

To obtain a random sample of events for the engine, we apply the following algorithm.

1. Get the number of trials (T) and number of years (N) in the experimental horizon for each trial from the user.

2. Assign a randomly generated "scramble seed" to each distinct year index in the database. The scramble seed will be identical for all events in same year, so that events within a simulated season are kept together and in order[14].

3. Scramble the set of simulated seasons (like shuffling a deck of cards) by sorting the data by scramble seed.

4. Determine the probability of an event-free year by dividing the number of distinct year indices in the database by the total number of years simulated by the cat model.

5. For each trial and each year within each trial, pick a random number between zero and one to determine whether the simulated year is storm-free or not (less than the empirical probability means event-free).

6. If the year is not event-free, assign a season's experience from the randomized database, starting at the top and moving downward (like dealing cards from the deck to years with events).

To build T observations of retained and reinsurance costs over each of N years, the engine does the following:

[14] Our cat model vendor informs us that each simulated event is independent. That is, both the frequency of storms from year to year and the severity of successive storms within a season are not related. But consider a next-generation cat model incorporating an "El Niño" module, whereby storm severity peaks during certain years, or a "Cool Current" module, whereby big storms tend to be followed by smaller storms if they take the same path (due to the thermodynamics of the ocean). In that case, keeping storms together and in order within the year would be vital to preserving the scientific parameters of the model. We saw little additional benefit in potentially disturbing the scientific parameters of the model for the sake of marginally better randomizaton.

1. Begin at the "top" of the simulated company experience, including the event-free years in their proper places.

2. Run the experience through the engine, by trial and by year within trial, to execute the financial model of FHCF incurred losses, cash, debt, and ultimately coverage afforded to the primary company.

3. Going back to the top, apply each season's event experience and its associated FHCF results to the private reinsurance program structure, then use the results of step 2 to integrate the public recovery, determining net costs to all parties for each season within each trial.

4. Output the simulated costs in each cell, as well as overall net retained costs and the net gain or loss for each reinsurance layer, to a database-ready flat file for export.

After the simulation logic in the spreadsheet does its work on the imported ELCID database, a post-processing step remains:

5. Use a desktop database application to assemble the observations of annual retained costs, thereby tabulating the empirical (simulated) distributions of C(1), ..., C(N). Use the distribution of retained costs and sorting functions to determine the empirical probabilities of C(i) exceeding the critical value KS in equation (1). Determine the expected values of C(i) with standard database queries.

We need to be more specific regarding step 3 above, namely determining the true costs to all parties in each layer for each season. The random variables of greatest interest to the ceding company are the retained costs in each layer. In the bottom (private program retention per event) layer, all losses are retained – this is R_r. In the top layer, the "over the top" or excess losses per event (R_x) are retained from the point of view of the private program, but can be subject to reimbursement from the public program. In the reinsured layers, the accounting equation for retained losses incorporating both the public and private programs is more complex:

$$\boldsymbol{R_c} = \boldsymbol{L} - \boldsymbol{V} + \boldsymbol{\pi_R} - \boldsymbol{B_R} - \boldsymbol{B_g} \qquad \textbf{(10)}$$

where R_c is the retained loss in the layer, L is the loss subject to the private program in the layer, V is the loss covered by the private program in the layer (reflecting the actual coverage levels selected), π_R is the reinstatement premium payable to private reinsurers in the layer, B_R is the portion of FHCF recovery used to refund reinstatements, and B_g is the portion of FHCF recovery used to fund "gap" losses. (Subscripts for storm and layer indexing have been suppressed.) Exhibit 4 diagrams the loss components of the private program before and after the introduction of public coverage (at, hypothetically, the 45% level).

Recalling §[5], in our experiment we allocate costs and recoveries in the following manner:

1. Calculate subject losses L_{ij} from each storm in each layer (including basic retention and over-the-top).

2. Calculate private program recoveries V_{ij} and reinstatement premiums due π_{Rij} for each storm in each layer.

3. Aggregate the preceding calculations to obtain season aggregate subject losses, private covered losses, and reinstatement amounts due in each layer.

4. Calculate the actual season aggregate recovery from the FHCF (B), reflecting the impact of any coverage failures determined by the dynamic model of the public facility.

5. Allocate the FHCF "pot" of recovered losses to (in order): over-the-top losses (B_x), reimbursement for reinstatement premiums in each layer (B_R), and "gap" losses (the portion of the losses between 95% and actual private coverage level selected) in each layer (B_g).

6. Since the basic retention and co-participation of 5% were specified in the private contracts, these losses are not reimbursed. The remainder of the FHCF money (B_v) is allocated to each layer's private reinsurers pro-rata to the amount of covered losses in each layer.

For the sake of completeness, we can view the simulation results from the point of view of the private reinsurers with a mirror-image equation:

$$V' = V - \pi_R - B_v \tag{11}$$

Net covered losses (V') equal contract covered losses, less collected reinstatements, less the reinsurers' share of the FHCF recovery. Note that a powerful by-product of the dynamic simulation is the ability to estimate the expected cost component of the private market reinsurance prices in each layer, as well as the variability of reinsured loss experience in each layer, which should be closely related to the cost of capital component of the price [Kreps 1990]. This is true whether or not the FHCF is incorporated into the experimental program.

The engine tracks every variable in equations (10) and (11) in order to complete step 4, output of the net loss experience for the company and all of its reinsurers to a database for distribution analysis.

[9] Summarizing Results and Making Decisions

The simulation exercise generates an enormous amount of data. Practically speaking, how do we boil down the results of thousands of trials for every prospective season to a one-page exhibit useful for the CEO in making the annual reinsurance buying decision? This is the purpose of step 5 in §[8] above.

The engine outputs records identified by the scenario name being tested, the trial number, season number, and a "record type" indicating the random variable (L, V, π_R, B_R, B_g, B_v, or net experience by layer R_c and V') being tracked for the given trial and season. A field is stored for each layer of the program, including retention and excess. With these keys, the output is copied into a friendly database format. Observing all trials for a given season, averages and extreme values can be calculated for any record type and layer – the empirical distributions of each random variable are embodied within the data set. The extreme values are found by virtue of relative frequency (empirical percentile) analysis.

Information on expected and extreme values of total costs is synthesized into an "accounting" exhibit using equations (10) and (11). Averages for each cost element are tracked for each layer and in total, and overall net experience is stated both at its average value and at particular return periods (the reciprocals of the associated percentiles). For example, the net retained losses in the "100-year season" are calculated by finding the trial which represents the 99th percentile of the observations of this variable generated by the engine.[15] The averages and extreme values for total costs C(k) follow with the addition of reinsurance premiums paid to private and public sources, since the only stochastic element of C is the retained losses R.

Exhibit 5 outlines how this type of presentation can be used to compare an integrated public/private reinsurance program with the current program for A-Florida. Multiple reinsurance scenarios generated "on the drawing board" can be tested by the simulation engine to determine the magnitude of expected cost savings, and simultaneously their vulnerability to failures in public coverage in later seasons of the decision horizon.[16] The shaded cells in Exhibit 5 represent stochastic values generated by the engine, and question marks represent parameters of the hypothetical alternative program or bottom-line results. While each scenario must be simulated with a separate run, the engine allows the same randomly generated storm set to be tested against many programs to determine the total costs and risk tolerance associated with each, because the user can disable randomization of the event set before each run. In this fashion, a schedule of scenarios, their expected costs by season, and their vulnerability functions (probability of

[15] In practice, we find it easier to present an average of the variable values for, say, the five records above and below the record marking the 99th percentile of total retained costs, to avoid spikes in any particular value which would distort the presentation.

[16] In these comparisons, the rate-on-line implied by the current market rates for 95% coverage in a private layer, or, say, 45% coverage from the FHCF, is assumed to be constant with respect to the co-participation selected. This means that the external reinsurance premiums in the total cost function are assumed to vary directly with the volume of coverage selected in each layer.

loss of a given percent of surplus loss over the horizon) can be examined to select the best integrated program among those tested.

Empirically, the "extreme values" table often demonstrates how different the CEO's risk tolerance constraint looks when phrased in terms of season experience instead of PML. For example, consider a simulation of 2,000 trials of a three-season experience period. The "100-year PARL" (probable aggregate retained losses) over the horizon would be given by sorting the records and finding the trial with the 20[th] worst value of (R_r + R_c + R_x) over three seasons. In our experience, the aggregate result in the specific years which yield the PMLs used in common catastrophe reports is practically uncorrelated to that for the years which yield the same percentile in PARL, at least until the return periods become very remote.

In boardroom English, the "100-year storm" and the "100-year season" are not the same thing – the "100-year season" might be the one with four Category 2 hurricanes each exhausting the basic retention but leaving the top-layer reinsurers untouched. On the other hand, the 250-year season may be the one with the "perfect storm" which blows Miami off the map. Depending upon the stated goals of the program and the exposure profile of the insurer, using reinstatement provisions to increase the "depth" of the program may be a better value than buying additional layers to handle severe events.

[10] Generalized Applications

Insurers not operating in Florida personal lines may wonder if our example is too narrow. The theory, the problem definition, and variations of the tools we have developed in this paper are applicable to a variety of catastrophe management problems involving multiple and possibly non-interacting funding sources. We shared this example for several reasons. First, primary carriers must develop the ability to integrate government-backed capital sources into marketing and reinsurance strategy. In Florida and several other states, the public sector has a hand in the insurance business. Future Federal legislation may enhance the role of government capital. Most government programs are likely to have some of the peculiar attributes we described. Second, the public nature of the FHCF means a wealth of financial information is available, making it a good candidate for an experimental capital integration problem. The example is important because the utility of the techniques we have developed is not limited to the Florida market or public sector capital. In the following paragraphs we present three examples of how our methods can be generalized to situations commonly faced by insurers underwriting risks subject to catastrophic loss.

First, consider an insurer using a funded cover for a portion of its catastrophe exposure. Funded covers typically grow during loss free periods and are occasionally depleted by events, exhibiting characteristics very similar to the FHCF. Insurers using this approach must design tools to integrate traditional reinsurance vehicles with their funded program and to handle an abrupt depletion of the funded cover.

Second, consider a primary insurer operating in a market with strict rate regulation and significant barriers to exit. If the insurer extensively uses reinsurance whose price is unregulated, it may face a sharp reinsurance rate increase following a major loss event and be unable to pass it on to its policyholders. Given a certain "budget" for reinsurance embedded in its primary rate levels, it may be unable to afford a portion of the coverage it previously bought. This coverage "failure" requires substitution of alternative (and, under our equilibrium presumption, more expensive) capital, with ensuing adjustments in underwriting and operational strategy. The insurer acting with foresight would analyze its reinsurance strategy over a multi-year horizon, incorporating a market pricing model, and perhaps secure reinsurance with multi-year pricing (even at a greater initial cost) or develop contingent sources of capital.[17]

Figure 2: Proportion of Loss By County

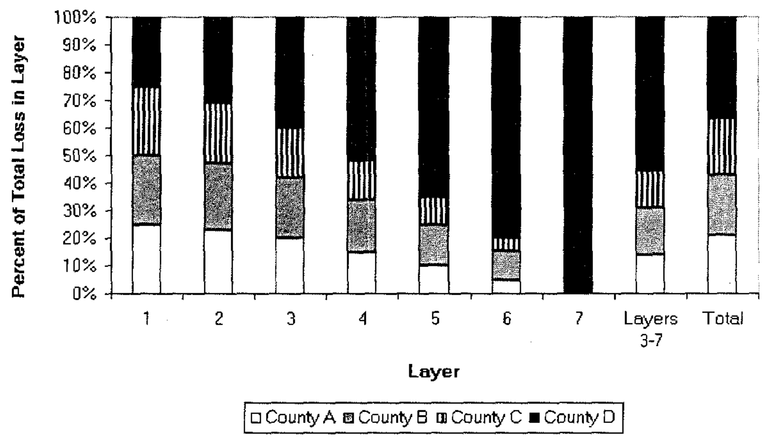

Third, consider allocating private reinsurance costs (either net or gross of FHCF recovery) to categories of policies using the event-level database. Figure 2 above shows a hypothetical distribution of losses by county calculated separately within each layer of season loss experience. Recalling loss distribution theory, the average volume of losses over a given number of observations always decreases in higher layers, while the

[17] In fact, this example is not unique to the insurance industry. California changed the structure of its electricity market in the late 1990's in such a way that the retail power providers (utilities) were disallowed from generating power internally and forced to buy power under only short-term contracts on an unregulated wholesale market, while the retail market maintained tight price controls. As retail demand grew rapidly and supply was restricted by external events, utilities faced steep price increases for capacity that could not be passed on to consumers, forcing them to the brink of bankruptcy in early 2001. The failure of public policymakers as well as the privately owned utilities to analyze and plan for this contingency was monumental.

coefficient of variation in upper layers is always higher. Immediate empirical implications are that the distribution of losses by county is not the same in every layer, and that external reinsurers charge higher risk loads (reflecting their increased cost of capital) in higher layers. These implications combine to form a significant reinsurance cost allocation problem in primary ratemaking – which counties deserve to bear most of the risk load?

A primary insurer purchasing catastrophe reinsurance for layers 3-7 might allocate the reinsurance cost to county using average annual losses by county from a cat model (the "total" column in Figure 2). This solution understates the needed premium in county D and, conversely, overstates premiums in counties A, B, and C. Allocating reinsurance costs based on the pertinent layers covered (the "layers 3-7" column in this case) would distribute the risk load to policyholders more equitably.

Even if prices are set with reinsurance costs allocated correctly, Figure 2 warns the insurer of the implications of writing new policies that may disturb the distribution of losses assumed in its pricing. In this example, the insurer can write additional policies in counties A, B, and C without buying additional layers of reinsurance. Alternatively, increasing county D exposure increases PML, requiring additional "vertical" reinsurance carrying a higher risk load than the average embodied in primary rates. At some point, marginal reinsurance premium on new business in county D could exceed marginal direct premium from the primary policies!

We believe that many enterprise-wide Dynamic Financial Analysis projects could benefit from incorporating these or other generalizations of the tools we have presented.

[11] Summary: Key Insights and Implications for Stakeholders

We conclude with a brief summary of the key concepts introduced in the paper:

1. The problem of selecting an optimal catastrophe reinsurance program can be expressed as a minimization problem for explicit costs (premiums and retained losses) with a risk tolerance constraint. This formulation is consistent with approaches in which the opportunity cost of capital is considered directly.

2. Public institutions can offer reinsurance coverage at below market prices by deferring part of the cost to the future through bonding and through favorable tax treatment. This cost of capital advantage is a result of the ability to break the accounting and tax rules governing the private sector, not an inherent advantage in portfolio composition or efficiency.

3. When assessing an insurer's exposure to catastrophic events, the traditional metric of gross single event "probable maximum loss" is inadequate. Only a standard considering "probable aggregate retained loss" over the decision horizon, with

event impact testing including the effects of both frequency and severity, is sufficient to capture the total loss distribution.

4. Both the FHCF and private reinsurers have volatile rates-on-line from season to season, but for different reasons. The FHCF charges constant premiums for a variable amount of coverage, while private reinsurers charge variable rates for a constant amount of coverage. FHCF and private ROLs oscillate harmonically, exacerbating market price cycles.

5. Because of the public sector cost advantage, the optimal reinsurance program for a Florida-only insurer involves successfully integrating private and public coverage. However, the näive strategy of replacing all possible private reinsurance coverage with subsidized public coverage exposes the primary insurer to a significant and measurable risk which must be compared to its risk tolerance.

6. The risk of fluctuations in public coverage can only be comprehensively analyzed by modeling both the cash position and debt burden of the FHCF over a multi-season horizon, which presumes development of industrywide event subject loss distributions.

7. The successful integration of season aggregate public coverage into a layered private reinsurance program requires careful definition of the inuring structure of the program to manage coverage overlap and coverage inequities. A simulation model, building on the dynamic analysis of the public reinsurer, can be used to model the integrated program structure and allocate costs, so that total cost distributions can be aggregated.

8. A model used to estimate a cedant's cost distributions can be applied equally well to estimate a reinsurer's cost distributions and check market prices.

9. Catastrophic events are the greatest single threat to the solvency of many insurers, but receive insufficient consideration in leading DFA models as well as other solvency testing tools.

10. Effective enterprise-wide DFA models must contain an ability to incorporate the net financial impacts of catastrophes, no matter how multifaceted the reinsurance structures of the company. This is true whether or not the disjoint capital sources are public or private. Our methodology is highly transferable to the DFA platforms used by many companies, and several parameters in our model could be greatly enhanced with the power of other parts of the DFA engine.

11. It is in the best interest of the public and its policymakers for each property insurer to optimize its individual protection against catastrophic events, particularly when significant reliance is placed on government-backed sources of capital.

Appendix A

Why Does a Catastrophe-Prone State Have So Many Single-State Insurers?

After Andrew, and with the help of the fledgling cat models, many insurers prospectively analyzed their exposure to hurricanes for the first time. The studies revealed that the magnitude of risk on the books was both unexpectedly high (relative to measurements based on historical experience) and unacceptable given the risk tolerances of management and the expected long-term return on the business. For policymakers, preventing the revelation from becoming a market collapse was paramount. Insurers were allowed to distribute risk in two new ways, by:

- Shifting the time horizon of their commitment to the Florida market.
- Shifting a portion of the expected policyholder deficit (over the revised time horizon) to the citizens of Florida.

These adjustments were accomplished by allowing the formation of single-state subsidiaries and affiliates.

Adjusting the length of commitment works because insurance profitability is volatile over restricted spaces and short periods, but predictable over a diversified book of business in the long run. A premise of this paper is that insurance company management objectives, explicitly or implicitly, include a time horizon over which a fair return is statistically assured within a certain tolerance. Continuing to write Florida business in the "main" company signals a long-term commitment to the market, because the profit or loss, regardless of magnitude, flows to the financial results of the main company – presumably a going concern - each year. The insurer knows that a fair return in Florida is assured only asymptotically[18], but no exit strategy is needed as long as the horizon is infinite. On the other hand, the profitability in a market which is cat-prone is relatively more volatile over short periods than that for one which is not. The time horizon required to assure a fair return (within the same tolerance) is longer. If management suddenly discovers that the length of commitment required in an evolving and uncertain market is unacceptable, the rational response would be to shorten the time needed for convergence as cheaply as possible.

One way to do so is to reduce exposure by non-renewing policies. This course of action is expensive due to existing fixed costs of doing business in the state, and in any event cannot be accomplished immediately. After Andrew, Florida regulators added an enormous additional cost to this strategy by requiring insurers to leave the state entirely or else adhere to strict limits on non-renewals of residential property policies.[19] A

[18] Inability to expect a fair return due to price suppression is another significant problem in the Florida property insurance market, but in this discussion we will assume rates allow a fair expected long-run return.

[19] The original 1993 moratorium, which continues to be extended and in effect today, prohibits an insurer from non-renewing more than 5% of its policies statewide or more than 10% in any county in a 12-month period.

palatable alternative to insurers, and one endorsed by government, was to allow an insurer to create a Florida-only subsidiary with a given initial capitalization, and renew all (and only) its property policies into the subsidiary. Using a gambling analogy, "playing the market" in Florida with a designated bankroll for a specific line of business limits the downside - it is economically identical to giving the insurer a free option to "quit the game" by folding up the subsidiary when it is in a deficit position of more than its original stake, while setting no restrictions on its gains. Without the subsidiary, the entire capital of the parent company would be supporting Florida catastrophe exposures. In another less volatile state, this risk may be tolerable compared to the certain transaction costs of setting up a subsidiary, but not so in Florida. The single-state company is thus an effective method of creating an option to shorten the commitment to the market.

A gambler can quit the table at any time, but his "chits" (casino credit) must be paid in full even if he has lost more than his original stake. In contrast, single-state subsidiaries have more than just the option to quit the state and pay their deficits by borrowing from the parent company. The existence of the Florida Insurance Guaranty Association (FIGA, the guaranty fund for the state) means the insurer also has the option to sell the eventual deficit of its subsidiary to the residents of Florida for the price of its initial capitalization. A general analysis of the value of this option can be found in Butsic [1994]. If the subsidiary is ruined by a catastrophe, the parent turns it over to FIGA and walks away. The guaranty mechanism in Florida, unlike that of many states, is ultimately funded by the residents. All remaining insurers in the state pay assessments to FIGA to fund the satisfaction of policyholder claims from the insolvent subsidiary, but the assessments are allowed by law to be recouped from all property-casualty policyholders in future premiums. It is not just the citizens who were insured by the insolvent subsidiary who are affected by the poor business decisions of a Florida-only company!

It should now be clear why an elected regulator agreed to the formation of Florida-only subsidiaries. Heuristically, we can decompose the needed premium to insure aggregate Florida property exposure over a period as ([Price per unit] X [Quantity] X [Time]). The regulatory strategy temporarily lowered the market-clearing price to a more politically palatable level (one net of the economic value of the gratis option), while fixing the quantity of insurance provided. In return, the time horizon was sacrificed; that is, part of the aggregate expected policyholder deficit was deferred for a while (with luck, at least until the next election) and post-funded by the citizens rather than pre-funded by the risk-bearers. The moral of the story? Frequently, the cost "savings" introduced by government modifications of private markets are actually just cost deferrals. When all economic costs are aggregated over a sufficiently long horizon, the mirage typically vanishes.

WORKS CITED

Bernstein, Peter L. Against the Gods: The Remarkable Story of Risk. New York: Wiley, 1996.

Musulin, Rade T. "Issues in the Regulatory Acceptance of Catastrophe Models for Ratemaking". Journal of Insurance Regulation, Spring 1997, p. 342.

Musulin, Rade T. "Taming the Big Cats". Best's Review, March 1999.

Kreps, Rodney E. "Reinsurer Risk Loads from Marginal Surplus Requirements". Proceedings of the Casualty Actuarial Society, 1990, p. 196.

State of Florida. Florida Statutes §215.555 and related adminstrative rules.

Meyers, Glenn G. "The Competitive Market Equilibrium Risk Load Formula for Catastrophe Ratemaking". Proceedings of the Casualty Actuarial Society, 1996, p. 563.

Florida Hurricane Catastrophe Fund. Annual Report, 2000.

Butsic, Robert P. "Solvency Measurement for Property-Liability Risk-Based Capital Applications". The Journal of Risk and Insurance, December 1994, p. 656.

Feldblum, Sholom. "Risk Loads for Insurers". Proceedings of the Casualty Actuarial Society, 1990, p. 160.

PRIVATE PROGRAM CURRENT STRUCTURE

	Program Layer						
	Retention	Working	1st Cat	2nd Cat	3rd Cat	4th Cat	Excess
Attachment Point	0	5,000,000	20,789,474	41,842,105	62,894,737	83,947,368	105,000,000
Exhaust Point	5,000,000	20,789,474	41,842,105	62,894,737	83,947,368	105,000,000	---
Layer Width	5,000,000	15,789,474	21,052,632	21,052,632	21,052,632	21,052,632	---
Source Covg. %	0%	95%	95%	95%	95%	95%	0%
Source Covg. $	0	15,000,000	20,000,000	20,000,000	20,000,000	20,000,000	0
Rate On Line	0.00%	20.00%	10.00%	6.00%	4.00%	2.00%	0.00%
Ceded Premium	---	3,000,000	2,000,000	1,200,000	800,000	400,000	---
Reinstatement Premium	---	3,000,000	2,000,000	1,200,000	800,000	400,000	---
Number of Covered Events	0	3	3	3	3	3	0

Estimated Subject Premium:	50,000,000	Retention as % of SMP:	10%
Max. Pro-Rata Recovery:	7,500,000	A-Florida Previous Year End Surplus:	75,000,000

PUBLIC PROGRAM FINANCIAL MODEL PARAMETERS

	"First" Season	"Second" Season
Max Incurred Loss:	11,000,000,000	11,000,000,000
Max Assessment:	641,400,000	320,700,000
Opening Cash Balance:	3,680,000,000	
Industry Annual Premiums:	445,000,000	
Industry Estd. Retention:	3,200,000,000	
Industry Avg. Coverage %:	85%	
Operating Expenses:	4,000,000	
Investment Yield:	6.00%	
Term of Bond Issues:	30	
YTM of Bond Issues:	6.04%	

A-FLORIDA PUBLIC PROGRAM OPTIONS

Coverage %	Retention	Premium	Coverage $
45%	15,000,000	1,000,000	25,000,000
75%	15,000,000	1,666,667	41,666,667
90%	15,000,000	2,000,000	50,000,000

Allowed LAE Load: 5%

FLORIDA HURRICANE CATASTROPHE FUND
TEMPLATE FOR DYNAMIC MODEL OF FINANCIAL POSITION

Exhibit 2

Item	Description	Formula
[1]	Industry Losses Subject to Public Coverage	given
[2]	FHCF Incurred Losses	([1] - [A]) x [B]
[3]	Begin Cash Balance	[C] in year 1 or [9] from prior year
[4]	Investment Income on Cash	[3] x [D]
[5]	Collected Premiums	[E]
[6]	Operating Expenses	[F]
[7]	Cash Available to Pay Losses	[3] + [4] + [5] + [6]
[8]	Paid Losses from Cash	Min ([2], [7], [G])
[9]	End Cash Balance	[7] - [8]
[10]	Total Debt Service on Prior Bonds	Sum all outstanding assessments for this season
[11]	Season 1	[H]
[12]	Season 2	If [10]=0, then 0, else [I]
[13]	Total Assessment Authority	Min ([11] + [12] , [10] + [11])
[14]	Remaining Assessment Authority	[13] - [10]
[15]	Theoretical Bonding Capacity	PV of an annuity of [14] at [J] and [K]
[16]	Actual (Capped) Bonding Capacity	Min ([15], [G] - [7])
[17]	Incurred Losses Excess of Cash Balance	[2] - [8]
[18]	New Bond Amount	Min ([16], [17])
[19]	Debt Service on New Bond	Annual payment to service debt of [18] at [J] and [K]
[20]	Maximum Industry Recoverable from FHCF	Min ([G], [2])
[21]	Actual Industry Recovery from FHCF	[8] + [18]
[22]	Industry Recovery Ratio	[21] / [20]; not needed if [20] is zero
[23]	Excess Industry Retained Losses	[2] - [22]

Item	Parameter
[A]	Industry Aggregate Retention
[B]	Industry Average Coverage Level
[C]	Begin Year 1 Cash Balance
[D]	Investment Yield
[E]	Industry Aggregate Ceded Premium
[F]	FHCF Operating Expenses
[G]	Max FHCF Incurred Losses
[H]	Max. First Season Assessment
[I]	Max. Second Season Assessment
[J]	Yield to Maturity of Debt
[K]	Term to Maturity of Debt

Exhibit 3

A-FLORIDA INSURANCE COMPANY
Raw Output from Event-Level Catastrophe Impact Database

Storm Information			Company Loss					Industry Loss
[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]
Year ID	Event ID	Order in Year	Gross Adjusted for Time Element	Pro-Rata Ceded	Capped Pro-Rata Ceded	Subject to Private Re	Subject to Public Re	Industry Subject to Public Re
1	3	1	5,792,630	754,336	754,336	5,038,294	5,094,397	1,143,484,928
2	4	1	157,535	22,949	22,949	134,586	92,476	125,765,088
5	6	1	15,677,770	2,303,284	2,303,284	13,374,486	11,662,850	3,846,696,448
6	8	1	11,511	481	481	11,030	9,847	209,397
7	12	1	445,828	58,255	58,255	387,573	382,144	65,985,676
9	15	1	1,460,098	43,044	43,044	1,417,054	1,392,059	25,274,374
9	16	2	9,238,030	234,963	234,963	9,003,067	8,861,356	158,270,048
9	17	3	10,183	393	393	9,790	8,815	179,405
11	20	1	763,519	116,046	116,046	647,473	606,143	184,306,864
12	22	1	4,517,121	616,976	616,976	3,900,145	3,806,402	951,313,024
12	23	2	5,817,446	199,385	199,385	5,618,061	5,531,140	71,945,944
14	24	1	60,754	3,396	3,396	57,358	56,862	1,690,202
15	25	1	27,788,411	1,595,003	1,595,003	26,193,408	25,794,787	991,666,240
15	26	2	29,242,270	2,588,661	2,588,661	26,653,609	26,235,477	1,649,879,552
15	28	3	3,282,911	419,087	419,087	2,863,824	3,019,619	393,474,208
16	30	1	1,696,668	114,661	114,661	1,582,007	1,484,225	103,848,360
18	35	1	6,379,298	816,664	816,664	5,562,634	5,698,554	1,044,632,640
18	36	2	796,594	19,501	19,501	777,093	761,725	14,041,427
19	37	1	19,278	431	431	18,847	16,200	280,985
19	40	2	7,808,022	389,584	389,584	7,418,438	7,342,077	265,341,360
20	41	1	4,395,080	629,330	629,330	3,765,750	3,216,568	959,188,800
20	42	2	2,057,726	301,027	301,027	1,756,699	1,611,333	459,089,568
24	45	1	9,673,032	421,290	421,290	9,251,742	9,232,142	217,995,776
25	46	1	2,589,420	63,518	63,518	2,525,902	2,484,141	41,941,136
25	47	2	58,812,042	4,034,229	4,034,229	54,777,813	51,591,372	3,574,883,584
25	49	3	1,256,312	179,212	179,212	1,077,100	1,151,592	237,913,760
28	51	1	15,315,979	872,980	872,980	14,442,999	14,128,880	736,751,744
29	52	1	6,589,651	883,393	883,393	5,706,258	5,886,333	1,012,107,648
29	54	2	56,122,492	6,614,285	6,614,285	49,508,207	47,146,442	6,823,095,296
30	56	1	2,079,291	308,664	308,664	1,770,627	1,735,815	457,605,344
30	58	2	1,387,237	53,988	53,988	1,333,249	1,326,659	35,287,520
31	59	1	528,140	76,992	76,992	451,148	431,778	117,309,544
31	62	2	244,710,448	36,501,169	7,500,000	237,210,448	202,759,252	46,213,873,664
31	63	3	18,763,264	598,427	598,427	18,164,837	17,632,150	368,258,816
32	64	1	237,895	12,987	12,987	224,908	210,021	10,359,245
34	66	1	40,918	2,386	2,386	38,532	36,704	1,341,627
37	68	1	600,619	14,975	14,975	585,644	572,949	10,759,818
38	69	1	65,048,736	3,701,529	3,701,529	61,347,207	57,592,892	2,166,251,520
38	71	2	15,409	464	464	14,945	12,990	242,705
45	77	1	920,804	39,340	39,340	881,464	872,286	19,535,520
46	79	1	106,019	8,640	8,640	97,379	78,220	9,383,332

Exhibit 4

ALLOCATION OF LOSSES IN ORIGINAL AND INTEGRATED A-FLORIDA REINSURANCE PROGRAMS

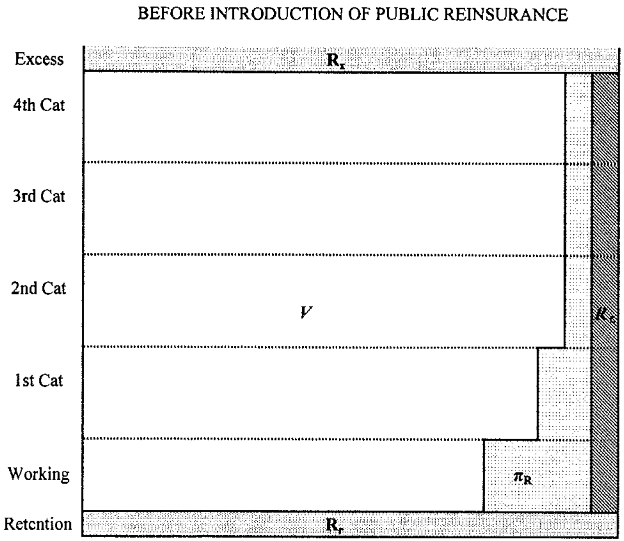

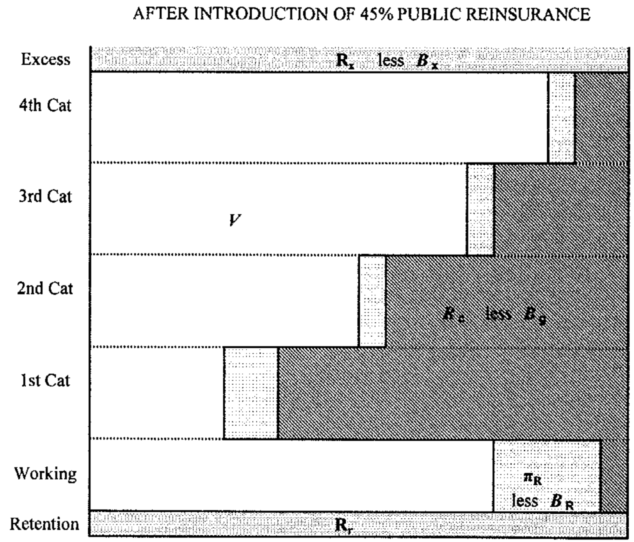

A-FLORIDA ACCOUNTING SUMMARY
FOR CURRENT AND ALTERNATIVE REINSURANCE OPTIONS

Exhibit 5

Simulation Engine Analysis (Expected Values)

		Private Program Layer Structure	All Layers	Retention	Working	1st Cat	2nd Cat	3rd Cat	4th Cat	Excess
[1]	Exhibit 1	Attachment Point		0	5,000,000	20,789,474	41,842,105	62,894,737	83,947,368	105,000,000
[2]	Exhibit 1	Exhaust Point		5,000,000	20,789,474	41,842,105	62,894,737	83,947,368	105,000,000	---
[3]	Exhibit 1	# of Covered Events		0	3	3	3	3	3	0
[4]	$E[L]$	Expd Private Subject Losses								
		Current Program								
[5]	Exhibit 1	Private Coverage %		0%	95%	95%	95%	95%	95%	0%
[6]	Exhibit 1	Private Coverage $	95,000,000	0	15,000,000	20,000,000	20,000,000	20,000,000	20,000,000	0
[7]	π_A	Private Ceded Premium	7,400,000	---	3,000,000	2,000,000	1,200,000	800,000	400,000	---
[8]	$E[V]$	Private Covered Losses								
[9]	$E[\pi_R]$	Reinstmt Premiums Paid								
[10]	$E[B_R]$	Public Recovery for Reinstmts								
[11]	$E[B_g]$	Public Recovery for Gap Loss								
[12]	$E[B_{.}]$	Public Recovery to Private Re								
[13]	$E[R]$	Net Layer Costs to Company	[13T]							
[14]	$E[V']$	Net Layer Costs to Private Re	[14T]							
		Alternative Program								
[15]		Private Coverage %		?	?	?	?	?	?	?
[16]		Private Coverage $		?	?	?	?	?	?	?
[17]	π_A	Private Ceded Premium	[17T]	?	?	?	?	?	?	?
[18]	$E[V]$	Private Covered Losses								
[19]	$E[\pi_R]$	Reinstatement Premiums Paid								
[20]	$E[B_R]$	Public Recovery for Reinsmts								
[21]	$E[B_g]$	Public Recovery for Gap Loss								
[22]	$E[B_{.}]$	Public Recovery to Private Re								
[23]	$E[R]$	Net Layer Costs to Company	[23T]							
[24]	$E[V']$	Net Layer Costs to Private Re	[24T]							

Total Cost Summary

	Public Covg $	[A] Public Re Premium	[B] Private Re Premium	[C] Net Loss Retained	[A]+[B]+[C] Expected Total Costs
Current Program	0	0	7,400,000	[13T]	
Alternative Program	?	?	[17T]	[23T]	
Difference					???

[13] = [4]-[8]+[9]-[10]-[11]
[14] [8]-[9]-[12]

[23] [4]-[18]+[19]-[20]-[21]
[24] [18]-[19]-[22]

Risk Tolerance Summary (Extreme Values)

	Net Loss Retained for Return Periods				
	10 years	20 years	50 years	100 years	250 years
Current Program					
Alternative Program					
Difference	???	???	???	???	???