Hedging Catastrophe Risk Using Index-Based Reinsurance Instruments

Lixin Zeng

Hedging Catastrophe Risk Using Index-Based Reinsurance Instruments

Lixin Zeng

March 2003

Abstract

Index-based hedging instruments such as industry loss warranties are increasingly recognized as effective hedging tools for insurance and reinsurance portfolios. However, wider adoption of these instruments is inhibited by basis risk, the difference between the index-based payoff and the buyer's actual loss. This study presents a systematic approach for potential buyers to analyze and manage basis risk in order to take full advantage of the benefits offered by these instruments.

We examine two measures of basis risk: (i) hedging effectiveness and (ii) conditional payoff shortfall. Many existing measures such as hedge volatility and correlation are special cases of the hedging effectiveness measure. Next, we study the tradeoff between basis risk and the cost of hedging. Finally, we present a robust numerical algorithm designed to optimize an index-based hedging program consisting of multiple index-based contracts.

1. Introduction

In recent years, we have observed growing interest in index-based hedging instruments, especially in the areas of catastrophe risk reinsurance and securitization. Examples include industry loss warranty (ILW) contracts and index-based cat bonds. In contrast to a traditional indemnity-based reinsurance contract, an index-based instrument has a payoff that is not completely determined by the loss incurred by the purchaser[1]. Instead, it is determined by an index that is positively correlated with the purchaser's actual loss. The index can be the industry loss or certain meteorological or seismic parameters related to a natural disaster event. The most frequently used industry loss indices used in the US are based on incurred insurance losses surveyed and published by the Property Claims Service.

The main advantage of index-based instruments is that they are practically free from moral hazard, a major hurdle that discourages capital market investors from participating in insurance risk securitization, even though the natural catastrophe risk is an extremely appealing asset class from a portfolio perspective (Litzenberger *et. al.*, 1996). The absence of moral hazard also suggests that an index-based instrument should command a lower margin than a comparable indemnity-based reinsurance contract (Cummings, *et. al.*, 2003), making it an attractive alternative to traditional reinsurance. Moreover, it is shown in Doherty and Richter (2002) that combining indemnity contracts with index-based instruments can ideally lead to efficiency gains for purchasers.

However, index-based instruments pose a new challenge to the purchasers in the form of basis risk - the difference between the actual loss experienced by the purchaser and the payoff of the index-based contract. The difference is one of the primary factors that have kept many potential purchasers away from these instruments. A systematic, credible, and practical way to quantify and manage basis risk must be made available to the potential

[1] Currently, the purchasers of index-based instruments are almost exclusively insurance and reinsurance companies. However, end users of insurance (e.g., corporations) have started exploring the use of this type of instruments.

purchasers before index-based instruments can gain recognition as a main stream risk management tool and become widely adopted.

The task of quantifying and managing basis risk can be divided into two problems: First, given an existing portfolio of liabilities to be hedged and an index-based hedging program consisting of one or more index-based contracts, how best to quantify the basis risk associated with this hedging strategy? Second, given an underlying portfolio and a set of constraints reflecting the buyer's risk appetite and return requirement, how can one construct an index-based hedging program to achieve an optimal balance between cost and hedging effectiveness?

This study focuses on these two issues. In Section 2, we state the assumptions and notations used in this study. Next, we develop an analytical framework to quantify basis risk in an effort to unify commonly used measures of basis risk (Section 3). In Section 4, we introduce an approach to construct an index-based hedging program that optimally balances hedging effectiveness and cost while satisfying certain constraints. Section 5 summarizes the study.

2. Assumptions and notations

We do not assume any specific form of parametric distribution for the random variables such as losses and underwriting profits. Instead, we represent the randomness of the "state of the world" using a large number of scenarios. This is because our primary interest is in hedging catastrophe risk and the outputs of most catastrophe models, which serve as inputs to our analyses, are scenario-based. In addition, although the numerical examples presented in this paper are realistic, they are hypothetical and are not based on any specific catastrophe model or actual company data.

Furthermore, we make three simplifications. First, it is assumed that only one loss event occurs in a year, although the analyses presented can be extended to include multiple events on an annual aggregate basis without difficulty using existing dynamic financial analysis (DFA) tools. However, not including DFA allows us to simplify the equations and focus on basis risk analysis. For the same reason, we also ignore premium reinstatement provisions frequently observed in actual transactions. Second, we do not consider the potential basis risk arising from the counterparty credit risk (i.e. the risk that the seller of the hedging contract fails to fully perform its contractual obligation). This permits us to focus on the discrepancy caused by the general lack of a one-to-one relationship between the actual loss and the index value. Third, we use binary ILW contracts in all examples. Nevertheless, the methodology developed can be applied to other forms of index-based instrument without substantial modification.

Lower and upper case letters are used to represent deterministic and random variables, respectively. Let L be the actual loss and X be the payoff of a hedging instrument. X is a function of an index I:

$$X = g(I) \tag{1}$$

For a binary ILW, I is the predefined industry loss for a region and given peril(s), and the payoff is defined as

$$X_I = g_I(I) = \begin{cases} l, I \geq i_t \\ 0, I < i_t \end{cases} \tag{1a}$$

where l is the limit of the ILW and i_t is known as the *trigger* of the contract. Another special case is an indemnity reinsurance policy, where $I = L$ and the payoff is defined as

$$X_R = g_R(L) = \begin{cases} l', L \geq r + l' \\ 0, L < r \\ L - r, r \leq L < l' \end{cases} \tag{1b}$$

where r and l' are the retention and limit of the reinsurance policy, respectively.

The net post-hedging loss L^* is then.

$$L^* = L - X = L - g(I) \tag{2}$$

It is possible that X > L. However, from an accounting point of view, this will not be allowed if the buyer wishes to treat the hedging instrument as reinsurance. Hence, Equation 2 frequently takes the following form

$$L^* = \max[0, L - X] = \max[0, L - g(I)] \tag{2a}$$

which forbids the buyer from claiming more than the actual loss. Specifically, we use L^*_I and L^*_R to denote the net loss after an ILW and an indemnity reinsurance policy, respectively:

$$\begin{aligned} L_I^* &= \max[0, L - g_I(I)] \\ L_R^* &= L - g_R(L) \end{aligned} \tag{2b}$$

3. Definition and quantification of basis risk

<u>3.1. The cause of basis risk – a qualitative view</u>

With an indemnity reinsurance policy, the amount of payoff is always precisely predictable given an actual loss, even though the actual loss itself is random (e.g., Equation 1b). However, this is generally not true for index-based instruments. We consider a hypothetical insurer (Company A), which has a geographically diversified exposure in the region where it sells property insurance and is considering using an ILW to hedge its catastrophe risk. As shown in Figure 1, at a given level of actual loss (e.g., along the dashed horizontal line), the industry loss index cannot be uniquely determined *a priori*. As a result, if Company A buys an ILW (Equation 1a) with a trigger represented by the vertical dashed line in Figure 1, the ILW payoff can be either zero or *l*, represented by the scenarios to the left and the right of the vertical line, respectively. This randomness makes it impossible for a buyer to precisely predict the payoff as a function of the actual loss. Next, we attempt to quantify such randomness, which is known as the "basis risk".

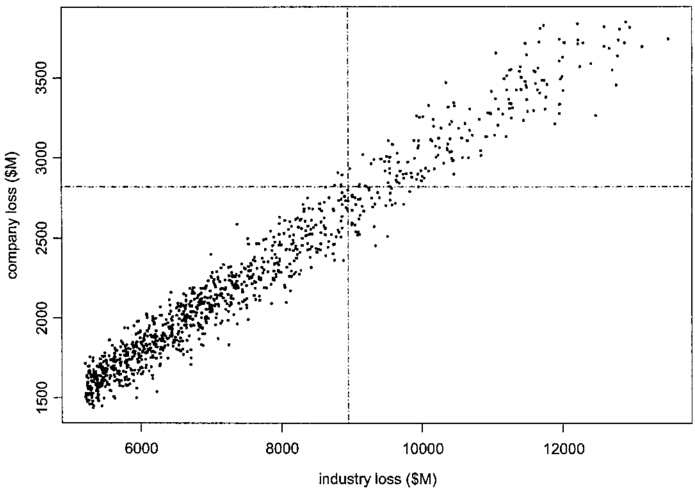

Figure 1. The loss to Company A vs. the industry loss. Each point in the plot represents a loss scenario. The dashed horizontal and vertical lines represent given levels of company and industry losses, respectively.

<u>3.2. Benchmarks for comparison</u>

Although basis risk is caused by the random difference between the index-based payoff (X_I) and actual loss (L), it is not sensible to directly compare X_I and L because rarely does a buyer expect the actual loss to be fully hedged. In the context of hedging catastrophe risk, the focus of the buyer is on reducing the severity of *large* losses. Hence, it is more meaningful to compare X_I to the payoff of a benchmark indemnity reinsurance policy (X_R) or, equivalently, compare the net losses associated with the index-based instruments and the benchmark, i.e. L^*_I vs. L^*_R (Cummings, *et. al.*, 2003).

The choice of the benchmark is usually based on the risk management objective of the buyer. For example, Company A currently has an annual probability of default[2] of 1%; a

[2] For illustration purpose here, the company is considered in default if the loss exceeds its surplus.

change in business environment requires this probability to be reduced to 0.4%. The traditional reinsurance approach to accomplish this is to purchase an indemnity reinsurance policy with the retention $r = v_0$ and the limit $l' = v_1 - v_0$, where v_1 and v_0 are the 99th and 99.6th percentile value at risk (VaR) of the underlying portfolio. Hence, its payoff (X_R) and net loss after this reinsurance (L^*_R) can serve as the respective benchmarks for the payoff (X_I) and net loss after an ILW (L^*_I). The cumulative distribution function of the loss of the underlying portfolio is shown in Panel I of Figure 2.

Next, we attempt to use an ILW to accomplish the same objective stated above. We choose the 99th percentile of the industry loss as the trigger and $v_1 - v_0$ as the limit, denoted i_t and l, respectively (Equation 1a). The basis risk of the ILW can then be defined based on the difference between L^*_I and L^*_R.

3.3. Definition and quantification of basis risk

The cumulative distribution functions (CDF) of L^*_R and L^*_I are shown in Panels II and III of Figure 2. Since L^*_R and L^*_I are random, we can compare their respective statistical summaries or evaluate the statistical summaries of their difference ($L^*_R - L^*_I$). These comparisons lead to the definitions of two types of basis risk.

Basis Risk Related to Hedging Effectiveness (Type I): In general, the purpose of purchasing a hedging instrument (reinsurance or ILW) is to reduce the risk of the underlying portfolio. The hedging effectiveness of the instrument can be measured by the amount of risk reduced. Let h_r and h_i denote the hedging effectiveness of the benchmark and the ILW. They can be defined as

$$\begin{aligned} h_r &\equiv 1 - y_r / y_g \\ h_i &\equiv 1 - y_i / y_g \end{aligned} \qquad (3a)$$

where y_g, y_r, and y_i are the statistical measures of the risk of the underlying portfolios before any hedging, net of the benchmark, and net of the ILW, respectively. Frequently used risk measures include standard deviation, value at risk (VaR), tail value at risk (TVaR), and probability of default (POD). The choice of the proper risk measure has been extensively discussed in the actuarial literature (e.g., Artzner *et. al.*, 1999) and is not repeated here.

The Type I basis risk (referred to as b_I hereafter) measures the hedging effectiveness of an index-based instrument relative to that of the benchmark. Hence, it can be defined as

$$b_1 \equiv 1 - h_i / h_r \tag{3b}$$

where we assume the benchmark hedging always reduces risk, i.e. $h_r > 0$.

Equation (3b) is obviously not the only valid definition. In fact, any b_I that increases with decreasing h_i/h_r is a valid quantification of basis risk. Partially due to this reason, basis risk is not uniquely defined in previous studies. For example, Major (1999) uses volatility of hedging to represent basis risk, whereas Harrington and Niehaus (1999) and Meyers (1996) measure basis risk based on the linear correlation coefficient between the actual loss and index-based payoff.

For Company A, the selected risk measure is the probability of default (POD), as reducing POD is its objective of hedging. Since the ILW does not reduce POD to the desired benchmark level, a substantial amount of basis risk exists (Table 1).

Table 1. Numerical values of hedging effectiveness and basis risk related to the ILW structure defined in Section 3 for Company A

	Underlying portfolio	Net of indemnity reinsurance	Net of ILW
Probability of default (risk measure)	1.00%	0.40%	0.60%
Hedging effectiveness		60.0%	40.0%
b_I			33.3%

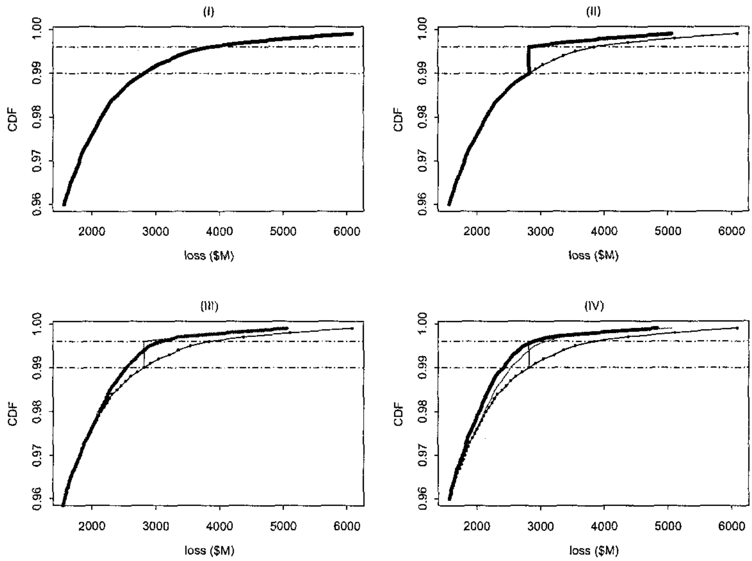

Figure 2. The cumulative distribution functions (CDF) of the gross and net losses of the underlying portfolio of Company A: (I) without hedge; (II) the thick curve: net of the benchmark; (III) the thick curve: net of the ILW defined in Section 3.2; (IV) the thick curve: net of the optimal ILW defined in Section 4.2. The thin curves in each of the panels (II), (III) and (IV) are the same curves as those in the previous panels for comparison purposes. The two horizontal dashed lines represent the 99% and 99.6% quantiles of the CDF.

Basis Risk of Payoff Shortfall (Type II): In general, two hedging instruments that accomplish the same hedging effectiveness do not guarantee the same payoff. Hence, even if b_I for an index-based hedging instrument is zero, it is still possible that the index-based payoff is less than the benchmark. To account for such discrepancy, we define the payoff differential (ΔL^*) as:

$$\Delta L^* \equiv X_I - X_R = L_R^* - L_I^* \tag{4}$$

where X_I, X_R, L^*_R, and L^*_I, defined in Section 2, are the index-based payoff, reinsurance payoff, net loss after the benchmark reinsurance, and net loss after the index-based product, respectively. A negative value of ΔL^* indicates that the buyer of the index-based instrument would recover more if the benchmark indemnity instrument were used instead (i.e. there is a payoff shortfall for the index-based instrument). This is another important aspect of basis risk in addition to its impact on hedging effectiveness. Because the purchaser is generally interested in protection against large losses, we examine the conditional cumulative distribution function of ΔL^* given the occurrence of a loss severe enough to trigger the payoff of the benchmark (i.e. $X_R>0$). The *conditional CDF* is simply denoted as $f_b(s)$:

$$f_b(s) \equiv prob(\Delta L^* < s \mid X_R > 0) \tag{5}$$

where $prob(\bullet)$ stands for the probability that "•" occurs. Examples of $f_b(s)$ are shown in Figure 3. Since we are primarily interested in measuring the downside risk of index-based instruments, we define the *Type II Basis Risk* (referred to as b_2 hereafter) as:

$$b_2 \equiv \frac{\max(-s^{\alpha},0)}{l'} \tag{6}$$

where s^{α} is the α^{th} quantile of $f_b(s)$. Under this definition, b_2 is the quantile of the index-based payoff shortfall normalized by the limit of the benchmark indemnity reinsurance policy (l'). For the ILW structure defined above for Company A, selected values of b_2 are listed in Table 2. The last row in the table shows that, for example, *given* the occurrence of a loss greater than the benchmark retention (r), there is a probability of 0.05 that the index-based payoff shortfall will exceed 19.9% of the limit of the benchmark hedging program.

Table 2. Selected values of b_2 for the initial ILW structure defined in Section 3.2 and the optimal ILW defined in Section 4.2 for Company A.

α	b_2 *(initial)*	b_2 *(optimal)*
0.004	43.4%	19.3%
0.01	41.1%	17.7%
0.05	19.9%	1.8%

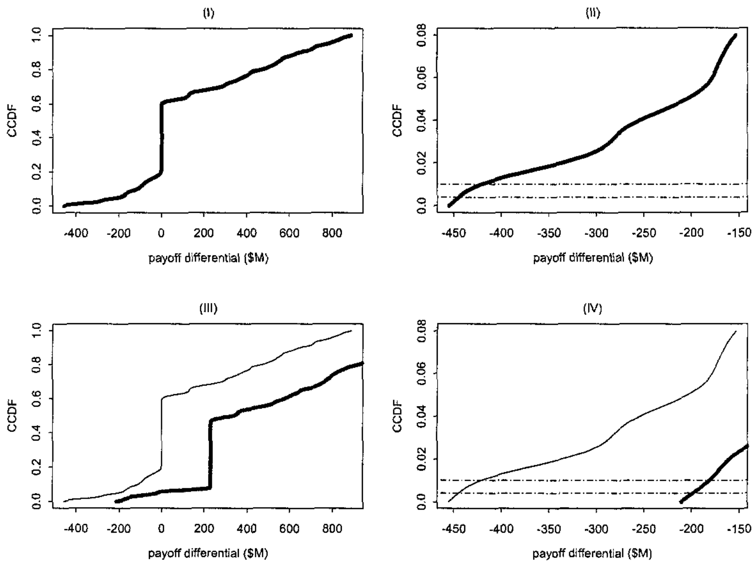

Figure 3. the conditional CDF of payoff differentials. Panels I and II: for the ILW structure defined in Section 3; the negative tail of the curve in I is a shown in II. Panels III and IV: for the ILW structure defined in Section 3.2 (thin lines) and for the optimal ILW defined in Section 4.2 (thick lines); the negative tail of the curve in III is a shown in IV. The horizontal dashed lines in Panels II and IV represent the 0.4% and 1% quantiles of the conditional CDF.

In summary, b_1 measures the hedging effectiveness of an index-based instrument relative to a benchmark, which is usually an indemnity reinsurance policy. Because this is directly related to the risk/return profile of the net post-hedge portfolio, we believe b_1 should be the focus of the buyer in evaluating the benefit of index-based strategies. However, b_2 is also important in practical decision-making as it measures the "probability of regret" for choosing an index-based instrument over a more traditional indemnity reinsurance policy. In this context, b_2 does not reflect or give any value to the fortuitous gain[3] available from the index-based instrument, which must be taken into account for the purpose of designing an optimal index-based hedging program (Section 4).

4. Optimizing an index-based hedging program

4.1. An overview

When the basis risk associated with an index-based instrument exceeds a tolerable threshold established by the purchaser, the contract terms must be modified such that the basis risk is reduced to the acceptable level. Given an underlying portfolio, there are primarily two ways to accomplish this: (a) changing the index or indices used by the contract and/or (b) modifying the parameters associated with each index (e.g., trigger and limit). It is possible that the cost of the contract will increase due to these changes. An optimal contract best balances the cost and benefit while satisfying the constraints imposed on the buyer. The process of arriving at such an optimal balance is illustrated using a simple example (Section 4.2). A robust method for optimizing complicated real world index-based contracts is introduced in Section 4.3.

4.2. A simple example

We revisit the example of Company A. We assume that the company wishes to reduce the basis risk associated with the initial ILW defined in Section 3.2 by changing the limit and trigger of the ILW. Specifically, it wishes to accomplish the following two objectives:

(a) Reduce b_l to zero (i.e. it requires that the ILW has the same level of hedging effectiveness as the benchmark). In this case, the task is to reduce the POD net of the ILW from 0.6% to 0.4%.

(b) Achieve Objective (a) with the lowest possible cost, allowing the underlying portfolio to retain the maximum possible net expected profit.

---

[3] the fortuitous gain is referred to as the excess recover from an index-based instrument relative to the benchmark (i.e. when $\Delta L^* > 0$). Under reinsurance accounting, it is impossible for the buyer to recover more than its gross pre-hedging loss.

Hence, the optimal ILW in this case is one that maximizes the net expected profit of the underlying portfolio while keeping POD from exceeding 0.4%.

We first plot how POD varies as a function of the ILW trigger and limit (the thick contours in Figure 4). The POD is represented by the contours of equal POD values. A point on a contour labeled x represents the trigger/limit combination of an ILW contract net of which the underlying portfolio has a POD of x. We call such a contour the *equal POD curve of* x (e.g., 0.4%). All points located to the upper-left of the curve correspond to POD less than x, and vice versa.

The initial ILW is represented by the solid square, which is located on the equal POD curve of 0.6%. For the POD to be reduced to 0.4% or less, the limit and trigger combination must be adjusted such that it is located on or to the upper-left of the equal POD curve of 0.4%. In fact, an ILW represented by any point on the equal-POD line of 0.4% can achieve the first objective.

We next examine the costs associated with different ILW contracts in order to accomplish the second objective. It is assumed that the premium for the contract is equal to five times the expected payoff, representing a typical profit margin of this type of contract in the market. With this assumption, the net expected profit[4] is calculated and visualized as the thin contours in Figure 4. A point on a contour labeled y represents the trigger/limit combination of an ILW contract, net of which the underlying portfolio has an expected profit of y. We simply call such a contour the *equal profit curve of* y. All points located to the upper-left of the curve correspond to net expected profits less than y, and vice versa.

[4] Defined as the premium of the underlying portfolio minus the sum of (i) the cost of the ILW, (ii) the expected value of the net loss, and (iii) other expenses. These quantities are formally defined in Section 4.3

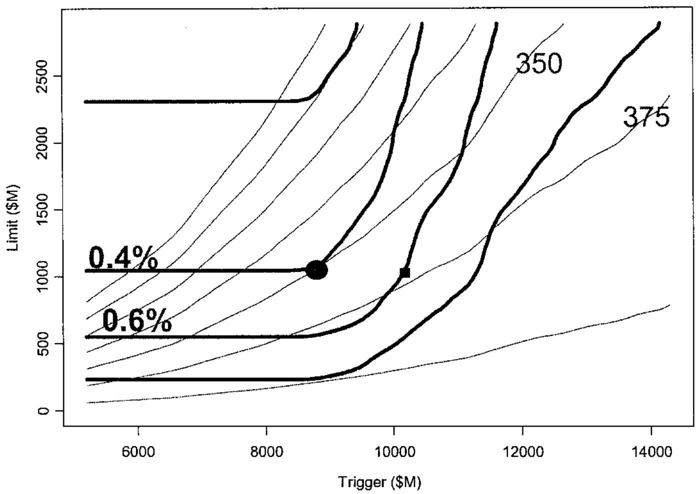

Figure 4. The probability of default (thick contour) and the expected profit (thin contour, in $M) net of ILW as a function of the trigger and limit. The solid square represents the initial ILW defined in Section 3.2, of which the index trigger is equal to the 100-year industry loss and the limit is equal to the difference between the buyer's 250-year loss and 100-year loss. The solid circle represents the ILW with the optimal trigger and limit arrived at in Section 4.2.

The point where the equal POD curve of 0.4% is tangent to an equal net profit curve is represented by the solid circle in Figure 4. The equal net profit curve represents a net expected profit of $350M. The solid circle represents the optimal ILW that accomplishes both objectives of the company because

(a) Since it is located on the equal POD curve of 0.4%, the first objective is achieved.

(b) All other points along and to the upper-left of the equal POD curve of 0.4% are also to the upper-left of the equal net profit curve of $350M. Hence, the net profits associated with these points are less than that associated with the solid circle. Thus, the solid circle represents the trigger/limit combination corresponding to the greatest net expected profit, i.e. the combination that accomplishes the second objective.

By definition, b_1 is reduced to zero. The loss distribution function of the underlying portfolio net of the optimal ILW is shown in panel IV of Figure 2. b_2 is shown in panel IV of Figure 3 and in Table 2.

This simple example shows that, in general, the task of optimizing an index-based hedging program is essentially a problem of optimally balancing basis risk and costs. Once the buyer determines the amount of acceptable basis risk and, if any, other constraints, an optimal hedging program should maximize an objective function specified by the user. In the example above, the objective function is the net expected profit. Other commonly used objective functions include risk-adjusted return on capital, Sharpe Ratio, etc. (e.g., Zeng, 2000). The optimization problem is formalized and generalized in the next subsection.

4.3. A robust method for optimizing an index-based hedging program

A robust method for optimizing an index-based hedging program is needed to handle real world tasks primarily because the underlying portfolio frequently consists of exposures in multiple lines of business and geographical regions. Thus, the number of indices involved is usually significantly greater than one. This makes the exhaustive search method used above impractical. In addition, it is not feasible to vary the limit and trigger continuously to create an ideal contract because only ILWs available in the market can be purchased. In fact, we can control only the amount to purchase for each contract.

For the k^{th} contract available in the market ($k = 1, 2, \ldots, m$), where m is the number of different contracts available, we define the following

I_k the underlying index (e.g. industry loss index for a specific region);

z_k the amount purchased;

$\eta_k(I_k)$ the unit payoff function.

p_k the unit premium (i.e. cost per z_k).

The payoff and cost of contract k are $z_k\eta_k(I_k)$ and $z_k p_k$, respectively. They are partitioned into the product of the amount of contract purchased and their respective unit values because the amount z_k is a decision to be made by the optimization procedure whereas the unit values depends on the contract itself, regardless the amount purchased[5]. For a simple binary ILW, p_k and z_k are simply the *rate on line* and the limit purchased, respectively. The payoff (Equation 1a) can be rewritten as

$$\begin{aligned} &\eta_k(I_k) = \begin{cases} 1, I_k \geq i_t \\ 0, I_k < i_t \end{cases} \\ &z_k = l \\ &g_l(I_k) = z_k \eta_k(I_k) \end{aligned} \tag{7}$$

The total payoff (X) and total cost (p_t) of the hedging program are

$$\begin{aligned} X &= \sum_{k=1}^{m} z_k \eta_k(I_k) \\ p_t &= \sum_{k=1}^{m} z_k p_k \end{aligned} \tag{8}$$

Hence, the loss net of the hedging program (L^*, defined in Equation 2b) can be specifically rewritten as

$$L^* = \max[0, L - \sum_{k=1}^{m} z_k \eta_k(I_k)] \tag{9}$$

The expected profit prior to hedging (EP) and the expected profit net of hedging (EP^*) can be expressed as

$$\begin{aligned} EP &= q_0 - EL \\ EP^* &= q_0 - p_t - EL^* \end{aligned} \tag{10}$$

[5] The unit premium actually depends on the amount purchased due to the supply-demand balance; however, this dependency is not considered in the analyses to simplify the formulas.

where q_0 is the inward premium of the underlying portfolio net of expense[6]. E is the expected value operator on a random variable.

The goal of the optimization procedure is to find the set of $z = \{z_1, z_2, \ldots, z_m\}$ such that a general objective function φ is maximized and a series of constraints are satisfied. Most frequently used φ is the expected profit of the underlying portfolio scaled by a risk measure. For the example, it can be defined as

$$\varphi = \frac{EP^*}{y_i} \tag{11}$$

where EP^* is the net expected net profit and y_i is some measure of the risk of the portfolio net of hedging. The latter can be the standard deviation, value at risk, tail value at risk and/or other statistics of the net loss L^*. The constraints can be expressed as

$$\begin{aligned} &\psi_c(EP, p_{i,} y_i, b_1, b_2) \le 0 \\ &c = 1,2,\ldots,n_c \end{aligned} \tag{12}$$

where n_c is the number of constraints. The constraints usually reflect limitations on the overall risk of the portfolio and/or the total cost of hedging. It is possible that a constraint can completely satisfy the risk control need of the hedger; consequently, the objective function does not need to be scaled by a risk measure, as illustrated in the simple examples in Section 4.2. In this example, there is one single constraint requiring that the probability of default net of hedging (denoted POD^*) do not exceed 0.4%. The objective is to maximize the net profit subject to this constraint. The objective and constraint for this example can be expressed as:

$$\begin{aligned} &\varphi = EP^* \\ &\psi_1 = b_1 = POD^* - 0.4\% \le 0 \end{aligned} \tag{13}$$

In general, given concrete expressions of φ, ψ_c, and y_i, which are chosen by the buyer of the hedging program, all the independent variables in Equations 11 and 12 are functions

[6] If L only contains catastrophe losses computed by a cat model, then the expected non-cat loss should also be excluded from q_0.

of z only. Therefore, the values of φ and ψ_c are functions of z only. Then, the optimization task becomes searching for z such that $\varphi=\varphi(z)$ is maximized, subject to $\psi_c=\psi_c(z) \leq 0$, $c=1,\ldots, n_c$.

If φ and ψ_c were linear or other smooth functions of z, this optimization task would be relatively easy to handle using traditional numerical algorithms such as the ones based on the steepest descent. However, because of the payoff function used in real-world transactions (e.g., Equation 7) are nonlinear and inherently not smooth, traditional optimization algorithms frequently fail to reach the global maximum.

In this study, an optimization procedure based on the genetic algorithm (GA) is used. Genetic algorithms are computing algorithms that simulate the mechanics of natural selection and natural genetics to "evolve" toward the optimal solution to problems. They are frequently applied to optimization problems where traditional approaches fail because of nonlinear, non-smooth, or discrete objective functions and constraints. A thorough discussion of GA is beyond the scope of this paper; however, interested readers can refer to, e.g., Goldberg (1989). The application of GA on index-based hedging is also introduced in Cummings, *et .al.* (2003). In this paper, we describe only the principle of this approach in the context of our task.

At first, randomly selected initial values are assigned to z to form the *original generation* (denoted z_0). Multiple individuals of the first generation ($z_{11}, z_{12},\ldots z_{1p}$) are created by randomly perturbing z_0, where p is the number of individuals; these p individuals are known as the *population* for this generation. A score for z_{1j}, based on the objective function and the constraint functions, is calculated to measure how "good" z_{1j} is. If any of the constraints is not satisfied, the score will be a large negative value (e.g., -10^{36}). If all constraints are satisfied, the score will be equal to $\varphi(z_{1j})$. The next generation population is created by combining two randomly selected individuals from the previous generation plus some random variations. The p individuals with the highest scores are retained (z_{21}, $z_{22},\ldots z_{2p}$). This process is repeated until a stopping condition is reached. For example, the stopping condition can be that the highest score among all populations in the current generation is very close to that in the previous generation. Upon stopping, the optimal z is the z_{gj} (i.e. the j^{th} individual in the g^{th} generation) with the highest score (hence the greatest value of the objective function) among all individuals. It is the random combination of individuals that allows the optimization procedure to "escape" from local maximums and have a much better chance to reach the global maximum.

We illustrate this approach using the following example. Company B has an underlying portfolio with exposure in two regions. It uses the 99th percentile VaR to measure risk (y_i); its goal is to reduce y_i to a target level while maximize the objective function defined by Equation 11. The objective and constraint are listed in Table 3. The ILW contracts available in the market are summarized in Table 4.

Table 3. Objective and constraint of optimizing an index-based hedging program

	Inward premium ($K)	Expected annual loss ($K)	Expected profit ($K)	99th percentile VaR ($K)	φ
underlying portfolio	10,000	2,305	7,695	54,861	14%
objective of hedging				less than 30,000	maximize

Table 4. Price and availability of ILW contracts

region	trigger ($M)	rate-on-line (p_k)	Capacity available ($M)	amount purchased ($\mathbf{z}$)
A	3,500	10%	20	z_1 = ?
A	10,000	6%	30	z_2 = ?
B	7,000	10%	25	z_3 = ?
B	20,000	6%	50	z_4 = ?

The task is to find the set of $z = \{z_1, z_2, z_3, z_4\}$ such that, net of the hedging program, the objective stated above is accomplished. In addition, the market data above impose another constraint: the maximum value of $\{z_1, z_2, z_3, z_4\}$ cannot exceed their respective available capacities (i.e. maximum limits). This example represents real world problems closely except a very small number of available contracts is used ($m = 4$), which allows us to verify the results using exhaustive search (i.e. testing all possible combinations of z_1, z_2, z_3, and z_4). Nevertheless, the speed performance of this approach for larger m is shown to be acceptable. Table 5 outlines the compositions of the hedging program recommended by GA and exhaustive search. The risk and return statistics of the portfolio net of the hedging programs are summarized in Table 6. Although the objective function considers only one risk measure (VaR), two additional commonly used risk measures (TVaR and standard deviation) are listed in the table for comparison.

Table 5. Optimal hedging program recommended by GA and exhaustive search. All values are in $K.

	z_1	z_2	z_3	z_4
Genetic algorithm	231	17222	24625	29563
Exhaustive search	0	17000	24500	29500

Table 6. Risk and return statistics of portfolio net of optimal hedging programs designed based on GA and exhaustive search. All values are in $K except the ratio φ

	Inward premium	Cost of hedging	Expected annual loss	Expected profit	99th percentile VaR	φ	99th percentile TVaR	Standard deviation
Underlying portfolio	10,000	-	2,305	7,695	54,861	14.0%	151,513	19,872
Net of optimal hedging – GA	10,000	5,270	1,312	3,419	14,419	23.7%	106,899	15,924
Net of optimal hedging – exhaustive search	10,000	5,240	1,317	3,443	14,641	23.5%	107,093	15,937

The GA-based results are very close to the benchmark solution produced by exhaustive search. In fact, it is better than the exhaustive search, in which the incremental value of z is only $500K. Although it is impossible to directly verify the results of the GA-based results using exhaustive search for larger m due to computational constraints, we believe that the GA-based algorithm remain accurate because it does not rely on any assumptions about m.

It is well known that most financial optimization procedures are subject to parameter risk, which can adversely affect the robustness of any optimal solution. For example, if TVaR is substituted for VaR as the risk measure in the objective function for the example of Section 4.3, the composition of the optimal hedging program will be different from that using VaR as the risk measure (Table 5). Generally, the solution can vary greatly depending on the choice of risk measure (e.g., VaR, TVaR or standard deviation), the parameter associated with the measures (e.g. percentile), or the mechanics of the underlying loss model (e.g., catastrophe model). For example, the optimal solution based on VaR would not necessarily be optimal if TVaR were used as the risk measure. Given the complex nature of this issue, using a coherent risk measure alone would not solve this problem. Although the parameter risk discussed above is not caused by or directly related to our optimization algorithm, the robustness of the outcome would be greatly improved if parameter risk could be handled more effectively, which remains a challenging problem for actuarial researchers and practitioners.

5. Summary

Index-based hedging instruments such as ILWs are increasingly recognized as effective hedging tools for insurance and reinsurance portfolios. However, wider adoption of these instruments is inhibited by basis risk, the random difference between the index-based payoff and the buyer's actual loss. This study presents a systematic approach for potential buyers to analyze and manage basis risk in order to take full advantage of the benefits offered by these instruments.

We examine two measures of basis risk: (i) hedging effectiveness and (ii) conditional payoff shortfall. Many existing measures such as the volatility of hedging (e.g., Major 1996) and R^2 (e.g., Harrington and Niehaus, 1999) are special cases of special cases of the hedging effectiveness measure, which quantifies the cost-adjusted benefit of the index-based hedging program relative to a benchmark. Next, we study the tradeoff between basis risk and the cost of hedging. The conditional payoff shortfall measures the probability that the buyer recovers less from an index-based hedging program than from a benchmark, reflecting the likelihood of "regret" for using the non-traditional hedging approach. In this study, a traditional catastrophe excess reinsurance layer is used as the benchmark. However, a wider spectrum of risk management products (such as proportional reinsurance, per risk excess reinsurance) is available. The methodology proposed in this paper is equally applicable to analyze these different benchmarks assuming the loss distribution of the underlying portfolio net of these products can be calculated.

Finally, we present a robust numerical algorithm designed to optimize an index-based hedging program consisting of multiple index-based contracts by analyzing the tradeoff between basis risk and cost of a hedging program. Compared to a benchmark optimization procedure based on exhaustive search, the GA-based approach is shown to work effectively to maximize the return on risk of a reinsurance portfolio subject to constraints. Nevertheless, like most financial optimization procedures, the outcome of this approach is not immune from parameter risk. Effectively address this issue remains a challenging but potentially rewarding future research direction.

Acknowledgement:

This paper would not have been possible without the encouragement, support and rigorous review by Bill Panning. The author would like to thank the CAS reviewers Tim Aman and David Drury for their extremely constructive input.

References:

Artzner, P., F. Delbaen, J.-M. Eber and D. Heath, 1999, Coherent Measures of Risk, Journal of Mathematical Finance, 9(3), pp. 203-28.

Cummins, J. D., D. Lalonde, and R. D. Phillips, 2003: The basis risk of catastrophic-loss index securities, to appear in the *Journal of Financial Economics*.

Doherty, N.A. and A. Richter, 2002: Moral hazard, basis risk, and gap insurance. *The Journal of Risk and Insurance*, **69**(1), 9-24.

Goldberg, D.E., 1989: Genetic Algorithms in Search, Optimization and Machine Learning, Addison-Wesley Pub Co, 412pp.

Harrington S. and G. Niehaus, 1999: Basis risk with PCS catastrophe insurance derivative contracts. *Journal of Risk and Insurance*, **66**(1), 49-82.

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, Special Issue Dec. 1996, 76-86.

Major, J.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).

Meyers, G.G., 1996: A buyer's guide for options and futures on a catastrophe index, Casualty Actuarial Society Discussion Paper Program, May, 273-296.

Zeng, L., 2000: On the basis risk of industry loss warranties, *The Journal of Risk Finance*, **1**(4) 27-32.