A Portfolio Management System for Catastrophe Property Liabilities

by Adam J. Berger, Ph.D., John M. Mulvey, Kevin Nish, and Robert Rush

Abstract

As catastrophe modeling systems become more sophisticated, the property insurance portfolio manager can receive better account loss information than ever before. We describe a software system called SmartWriter which effectively processes this information for the portfolio manager. Specifically, the system determines:

- Appropriate pricing for an account
- Which accounts to remove from a portfolio to maximize risk-adjusted return
- How to merge two books of business
- Where to grow or shrink business geographically to achieve maximum diversification benefits

We utilize a number of optimization techniques to address these issues. We formulate the problem as a large mathematical program with numerous loss scenarios (10,000 or more). We then describe an algorithm to solve the resulting stochastic optimization problem in order to maximize risk-adjusted return, expected utility, or other user-defined performance measures.

The SmartWriter system is a PC-based Windows application. USF&G, a large property and casualty insurance company, currently employs SmartWriter as an integral part of its decision making process.

1. Introduction

The insurance portfolio manager and underwriter require sophisticated analytical tools to assist decision making. Just as an asset portfolio manager, such as a mutual fund director, can immediately see the effects of adding a security or option to his portfolio's risk and return profile, the insurance portfolio manager needs to understand the effects of adding an additional account to the business line. In addition, there are many other issues the manager must address, such as: (1) Should an existing account be renewed and, if so, at what price? (2) Where are the best areas to expand the current portfolio? (3) How can two books of business be merged profitably?

We have developed a decision support system, called SmartWriter, which answers these questions for one application area, the catastrophe property business. SmartWriter employs data from earthquake and hurricane modeling systems to show the effects of adding a new account or subtracting an existing account from the current portfolio. In addition, SmartWriter optimizes the portfolio composition to produce a portfolio meeting user-specified characteristics. Although we are describing SmartWriter in the context of catastrophe property, the methodology applies to Directors & Officers, Errors & Omissions, Workers Compensation, and other insurance lines.

The paper proceeds as follows. Section 2 describes the method for evaluating an incremental account and the return on capital methodology. Section 3 lays out the optimization model to address the questions raised above. The algorithm for solving the problem is described in Section 4, and results are presented in section 5. We conclude with some next steps in Section 6.

2. Modeling an incremental account

Suppose we have a portfolio of insurance liabilities. As an example, we look at a portfolio of commercial businesses insured against earthquakes in California by USF&G, a large property and casualty insurance company. A potential new piece of business is presented to the portfolio manager, who must decide whether to write the account or reject it. Of course, some negotiating with the insurance broker who presents the account is possible, so the portfolio manager would also like to know the required premium to meet a profitability hurdle. Before analyzing the incremental business, we need to define a profitability measure for the existing portfolio. Two measures are return on allocated capital and expected utility.

2.1 Return on Capital

In this method, capital is assigned to a portfolio (or business unit) based on the risk of the portfolio. Risk is calculated based on characteristics of the cumulative loss distribution and portfolio profitability. For the catastrophe property business, capital is often a function of points in the tail of the distribution, similar to Value at Risk (VAR). For simplicity, we focus on a single point of the loss distribution, the 99^{th} percentile, and calculate the capital requirements as the funds needed to pay the loss incurred at the 99^{th} percentile. This is referred to as the "1-in-100 year loss", since one would expect the loss to be as bad or worse than this level once every hundred years. More complex formulas based on multiple points of interest on the loss distribution are possible (see Mulvey et al, 1997). Equation 1 shows the allocated capital calculation:

$$\text{capital} = \rho\, F^{-1}(0.99) - (p - e) \qquad (1)$$

where ρ is a discount factor, F is the cumulative distribution function for the loss, p is the premium received and e is the non-catastrophe expenses. The discount factor ρ is necessary since we receive premiums and pay out expenses (e.g. commissions) at the beginning of the year, and losses are incurred during the year. Thus, we calculate the capital required at the beginning of the year, and discount losses, so that all terms are on the same basis.

To calculate return, we first define *expected catastrophe loss* as the expected value of the loss distribution. Expected margin is simply premium less expense less expected cat loss. Expected *return on capital* (ROC) is calculated by dividing expected margin by the allocated capital:

$$ROC = (p - e - Ef(x)) / \text{capital} \qquad (2)$$

where f(x) is the loss distribution and Ef(x) is the expected value of the distribution.

We define the *marginal capital* for an account as the difference in capital required for the portfolio with the account and the portfolio without the account. Return on marginally allocated capital (ROMAC) has the same expected margin for the numerator and marginal capital in the denominator. Define c_p as the capital required for the portfolio with the account and $\underline{c}_p$ as the capital for the portfolio excluding the account. Then the marginal capital m_o and the return on marginal capital (ROMAC) is defined as:

$$m_o = c_p - \underline{c}_p \qquad (3)$$

$$ROMAC = (p - e - Ef(x)) / m_o \qquad (4)$$

ROMAC captures the diversification benefit of the account with respect to the portfolio. An account with a high ROMAC doesn't require much additional capital allocation for the portfolio as a whole, and thus is a good diversifier. Conversely, an account may have a high return on capital on a stand-alone basis, but a low ROMAC, and thus is most likely located in an area of high concentration.

To facilitate combining loss distributions, we discretize the sample space and create numerous scenarios. Each scenario represents a year's worth of catastrophes. We can then determine losses for the account in each scenario and combine accounts scenario by scenario to determine portfolio losses. Although it is not necessary to have scenarios for the above calculations (since capital with and without an account can be calculated separately with no need for combining accounts), it will be important in performing the optimization described in Section 3.

2.2 Expected Utility

An alternative approach to allocated capital is expected utility. Given an asset position for a business line (or company) at the start of a year, define a von Neumann-Moregnstern utility function over the range of possible asset positions at the end of the year (see Bell [1995] for an example). Each portfolio will then have an expected utility value calculated from its loss distribution. Portfolios can be compared simply on expected utility, with higher expected utility being more desirable. To see whether to add an account to a portfolio, compare the expected utility before and after the addition.

There are advantages and disadvantages to the return on capital and expected utility approaches. Return on capital is a familiar financial ratio and is easily explained. Allocating capital based on points on the loss distribution is straightforward and captures, to some extent, the risk inherent in the business. Unlike expected utility, however, the return on capital does not provide a definitive answer on whether to add a new piece of business (e.g., if a new account has below average return on capital and above average ROMAC). The expected utility framework takes into account all points of the loss distribution whereas return on capital methods generally incorporate only a few. Expected utility provides a definitive recommendation, but does not have an immediately intuitive interpretation. For example, a portfolio manager can appreciate that adding a new account will boost return on capital from 12.0% to 12.5%, but may not as readily interpret increasing expected utility from 3.45 to 3.47. Depending on the model, expected utility can be easier to solve (see Berger [1995]) since it fits more easily into a mathematical programming framework than return on capital, which requires sorting a discrete distribution; Sections 3 and 4 discuss this issue further. This comparison is summarized in Table 1.

	Advantages	Disadvantages
Allocated Capital	- Easy to explain - Returns have intuitive meaning	- Extra work to sort discrete distributions - Limited points on loss distribution
Expected Utility	- Handle entire loss distn. at once - Convex math program	- Hard to determine utility function - Results not intuitive

Table 1: Comparison of allocated capital and expected utility objective functions

2.3 Sample Decision

We present SmartWriter analysis (Table 2) of an account recently offered to USF&G's commercial property business. Although we have altered the raw output to protect confidentiality, the returns are consistent with the actual analysis.

	New Account	Current Portfolio	Combined
Premium	$980	$3,800	$4,780
Expenses	$294	$1,140	$1,434
Expected Catastrophe Loss	$71	$615	$686
Expected Profit	**$615**	**$2,045**	**$2,660**
Loss at 99th% = $F^{-1}(0.99)$	$5,200	$14,300	$18,100
Capital Required	$4,200	$11,600	$14,700
Return on Capital: ROC	**14.6%**	17.6%	18.1%
Ret. on Marginal: ROMAC	**19.8%**		

Table 2: New account analysis. All numbers in ($000), except where indicated.

The SmartWriter output is divided into three columns. The first column is the new account as a stand alone business. The expected income for the account, after taking expenses and expected catastrophe losses from the premium, is $615,000. The new account requires $4,200,000 in capital based on the 1-in-100 year loss of $5,200,000. This yields a return of 14.6%, which is below our hurdle rate of 15%.

The second column contains data on the portfolio as it stands today, and the final column is the portfolio performance if the new account were added. Note that the capital requirement for the combined portfolio is less than the sum of the new account and current portfolio capital: This indicates that the new account will diversify the business to some extent. Two additional items help quantify this diversification. The ROMAC for the new account is 19.8%, which means that the marginal return for adding the account divided by the marginal capital is significantly over the hurdle rate. The second item to note is the increase in the ROC for the portfolio from 17.6% to 18.1% if the account is added. For these reasons, the account was considered a good prospect, even though on a stand alone basis it was slightly below the hurdle rate.

3. Optimization Model

Optimization is the process of finding good solutions to a mathematical model. In the context of insurance underwriting, several problems are amenable to optimization. For a portfolio of large commercial accounts, the optimizer could locate the five accounts most in need of repricing, or the subset of the current portfolio which maximizes return. For a homeowners portfolio, the book of business is managed less on a home-by-home basis and more on a zip code, county, or state level; the optimizer can focus on which counties to expand market penetration and which zip codes to reduce premium volume. The next section describes SmartWriter optimization for commercial portfolios, and the following section for homeowner books.

3.1 Commercial Portfolios

Section 2 defined a method for comparing portfolios of accounts, either with return on capital or expected utility. We can now formulate an optimization model for choosing an optimal subset of accounts for the given portfolio. As mentioned above, we will define a discrete set of scenarios, where each scenario represents a number of catastrophes for a year. This facilitates the problem of combining loss distributions. For general continuous loss distributions, there is no simple method that can be used.

3.1.1 Variables and Objective

Define the following sets:

$\{1, 2, \ldots N\}$ – set of accounts in the portfolio
$\{1, 2, \ldots S\}$ – set of loss scenarios

Define the following input parameters:

p_i = premium for account i
e_i = non-catstrophe expense for account i
l_{is} = loss (in dollars) for account i in scenario s
π_s = probability of scenario s
ρ = discount factor

Define the following decision variables:

x_i, $i=1,\ldots,N$ – amount of account i in the portfolio

Our objective is to maximize return on capital:

$$\text{Max } \Sigma_{s=1,S} \Sigma_{i=1,N} \; \pi_s (x_i (p_i - e_i - l_{is})) / [\rho \, F^{-1}(0.99) - \text{sum } x_i (p_i - e_i)] \quad (5)$$

where $F^{-1}(0.99)$ is calculated from the revised loss distribution $x_i{*}l_{is}$.

Note that correlations are implicitly captured in the analysis. Since the entire loss distribution is calculated for the objective function, the correlation among accounts will affect the return on capital.

3.1.2 Constraints

The following are constraints that can be added to the model.
An account can either be in the portfolio or out of the portfolio so we add a binary constraint

$$x_i \in \{0,1\}$$

If one or more properties must be retained, we add:

$$x_i = 1$$

The total premium for the portfolio can not be reduced past a specified level, MinPrem:

$$\Sigma_{i=1,N}(x_i * p_i) >= \text{MinPrem}$$

The expected income on the portfolio can not be reduced past a specified level, MinInc:

$$\Sigma_{i=1,N} (x_i *(p_i - e_i - l_{is}) > \text{MinInc}$$

3.2 Expansion problem

Another problem facing insurers is where to grow a portfolio of a large number of small accounts, for example the homeowners market in California. These portfolios can not be analyzed account by account, since underwriters do not have the flexibility of choosing to write one home and not another. Accounts must be aggregated to a meaningful level: not too large so that accounts within a group possess similar characteristics, but not too small so that they can be managed effectively, such as with target marketing. The following model chose the zip code level as a reasonable trade-off between these competing demands. The objective function remains the same, but we change a few variable and constraint definitions. Our emphasis now is determining how much premium to retrieve from each zip code. We assume that the loss characteristics within a zip code are constant. Zip codes where this is not the case can be broken down into smaller units.

Define the following sets:

{1, 2, ... Z} – set of zip codes in the region
{1, 2, ... S} – set of loss scenarios

Define the following input parameters:

e = non-catstrophe expense ratio
l_{zs} = loss per dollar of premium in zip code z in scenario s
π_s = probability of scenario s
ρ = discount factor

Define the following decision variables:

x_z, $z=1,\ldots,Z$ – amount of premium from zip code z in the portfolio

Our objective is to maximize return on capital:

$$\text{Max } \Sigma_{s=1,S} \Sigma_{z=1,Z} \ \pi_s (x_z - e^* x_z - l_{zs}{}^* x_z)) / [\rho F^{-1}(0.99) - \Sigma_{z=1,Z} (x_z - e^* x_z)]$$

where $F^{-1}(0.99)$ is calculated from the revised loss distribution $l_{zs}{}^* x_z$.

Constraints similar to the ones in the pruning example above can be added; we give a few examples here. The premium level across zip codes can be bounded between two values, MinPrem and MaxPrem:

$$\text{MinPrem} < x_z < \text{MaxPrem}$$

Alternatively, the total expansion of the portfolio can be limited to a dollar value, MaxPort:

$$\Sigma_{z=1,Z} x_z <= \text{MaxPort}$$

4. Solution Procedure

The models described in the previous section are not easily solved with traditional mathematical programming procedures, due to the necessity of the sorting during the capital allocation calculation. We employ a number of metaheuristic search procedures to find the global optimum value for the problem. For all of these, it is important to find good starting points, which we describe first, followed by the search algorithm.

4.1 Elite Solutions

Elite solutions are points in the decision space which are believed to be good locations for a local search (also called intensification, since the local area is being explored thoroughly). One method for generating elite solutions for this example depends on the profitability of the portfolio as a whole and on the individual accounts. If the portfolio is profitable, then a candidate elite solution would be the entire portfolio, or the portfolio with a small subset of poor performing accounts removed. Alternatively, for a poorly performing portfolio, a candidate elite solution could be a small subset of profitable accounts, or no accounts at all. Another approach ranks accounts by profitability and correlation with the portfolio as a whole; an account with high profitability and low correlation would be included in an elite portfolio.

A more profitable approach relies on problem-specific information. Suppose the optimization procedure is run monthly or quarterly. Optimal solutions from previous runs can be stored and will provide good elite solutions, even if the portfolio has changed measurably since the last run.

Of course, accounts no longer in the portfolio but in the previous optimal solution must be removed.

After a number of elite solutions have been generated using some or all of the methods above, the solutions are ranked in terms of attractiveness. This ranking will then determine the order for the local searches (see next section). Ranking can be based on objective function value alone, but to fully explore promising areas of the decision space we can use a weighted average of the objective function and the distance from higher ranked elite solutions. As more solutions are ranked, the benefit for diversification increases.

4.2 Tabu Search

Tabu search was originally developed by Glover and has proven highly effective for solving combinatorial optimization problems. (See Glover [1989] for an introduction). The procedure searches a feasible region by monitoring key attributes of the points that comprise the search history. Potential search iterates possessing attributes that are undesirable with respect to those already visited become tabu; appropriate penalties discourage the search from visiting them. We provide details below.

Consider a general non-convex optimization problem of the form:

$$\underset{x}{\text{minimize}}\ f(x),\ x \in X$$

where the function f(x) corresponds to the return on capital objective in Equation 5.

Our adaptation of tabu search has three basic elements:

- a function $g(x) = f(x) + d(x) + t(x)$. The function $d(x)$ penalizes x for infeasibility. The function $t(x)$ penalizes x for being labeled tabu.
- the current iterate x_c,
- a neighborhood of the current point N_c.

The procedure generates a new iterate x_{new} by selecting the element of N_c for which g(x) is smallest. The tabu restrictions represented in t(x), can address short-, intermediate-, and long-term components of the search history. Short-term monitoring is designed to prevent the search from returning to recently visited points, allowing the procedure to "climb out of valleys" associated with local minima. Short-term monitoring can also serve as a rudimentary diversification vehicle. Intermediate- and long-term monitoring techniques provide for a much more effective diversification of search over the feasible region. In addition, the elite solutions described previously also provide diversification. See Glover [1990] for additional details.

Details of four processes are required to define our adaptation of tabu search: formation of the neighborhood of the current point, assignment of tabu penalties, termination of search procedure, and *greedy selection* of the new iterate from the neighborhood of potential moves.

Neighborhood formation proceeds as follows. Let $x_c = (x_{1c}, \ldots, x_{nc})$ be the current point; the decision vector thus has n components. For the example in Section 3.1, this would be a vector of zeros and ones, where a "one" indicates the account is in the portfolio. Each member of the neighborhood of x_c, N_c, is formed by modifying one of its components either up or down by an amount equal to some value *step_size*. Note that this operation implicitly defines a discretization of the continuous feasible region. There are thus 2n members of N_c. We call each of these members a potential move; one of these will become the new iterate, i.e. - the actual move. Each potential move is characterized by two move attributes: *index changed* and *new value*. Attribute *index changed* is equal to j, where x_{jc} is the component of x_c whose value is changed by the potential move; *new value* is the value that the component being changed by the potential move assumes (formally: *new value* = x_{jc}, such that j = *index changed*).

The manner in which we assign tabu penalties -- and thus define the function t(x) -- to each potential move relies on exploitation of short-term search history; the methodology is based on the technique developed in Glover, Mulvey, and Hoyland [1996]. The assignment is based on a comparison of the move attributes of each potential move and those of the iterates comprising the recent search history. The maintenance of two data structures is necessary: 1) the *tabu list*, and 2) *time of last change list*. The *tabu list* is composed of the attributes of the T most recent search iterates; *tabu list* is thus a T x 2 array where T = TABU LIST SIZE. The *time of last change list* is an n x 1 array, where *time of last change list*$_j$ = the last iteration during which the actual move's *index changed* attribute equaled j. We also define f_{BEST} as the best objective value (in terms of minimization) found by the procedure at any point in the search process.

Three criteria govern our assessment of the tabu status of each potential move (x_p):

Condition 1: do the move attributes of x_p match any of those in the *tabu list*?
Condition 2: is length of stay < REQUIRED STAY, where length of stay
= current iteration - *time of last change*$_j$,where j = *index changed*?
Condition 3: is $f(x_p) < f_{BEST}$ <u>and</u> $x_p \in X$?

If either of the first two Conditions are true, we assign an appropriate tabu penalty to the potential move, discouraging the search from moving to x_p. Condition 1 prevents the search from revisiting a point whose move attributes match those of points recently visited. It is this operation that allows the search process to move away from local minima, as we described earlier. Condition 2 insures that a variable is not changed too soon after it becomes the basis for an actual move; it thus is a vehicle for short-term search intensification If the final condition is satisfied, we eliminate the tabu penalty for x_p; this allows the search to move to a tabu point if the objective value associated with this point is better than that of the best point found thus far. (This is our implementation of the concept of *aspiration criteria*; we refer the reader to Glover [1990] for details.)

We present three termination criteria:

1) Total time exceeds a preset maximum
2) Total iterations exceed a preset maximum
3) The amount of time spent without any improvement in the solution exceeds a preset maximum

Finally we address the greedy criterion for selecting from the set of potential moves the actual move, and thus the new iterate. The standard approach for selecting the new iterate is to find the point in the neighborhood of the current iterate for which g(x) = f(x) + d(x) + t(x) is smallest, a process that by definition requires evaluation of every member of the neighborhood. This strategy can degrade the effectiveness of the search when the computational effort required to evaluate f(x) is prohibitive. The greedy search strategy addresses this difficulty. It calls for the evaluation of the set of potential moves to cease when a neighbor fax is found which $f(x) < f(x_c)$ and $d(x)= t(x)= 0$, i.e. – x is feasible, not tabu, and shows improvement.

5. Results

Below is the SmartWriter output for a California earthquake portfolio with 173 accounts. The results are from real company data, but the numbers have been disguised to protect client confidentiality. We ran the analysis on a Windows 95-based PC with 64MB of memory, with run time between 5 and 10 minutes, depending on parameter settings.

The optimizer recommended the removal of 16 accounts from the portfolio. Table 3 shows summary information before and after the optimization for the portfolio as a whole.

On the whole, this was a profitable book of business, but there were a small number of poorly performing accounts. Not only did these accounts have a poor expected return, but they had a severe effect in the tail of the distribution. Expected income only decreased by $100,000 (3%), but the loss at the 99th percentile decreased by over $15MM. Return on capital jumped from 14.7% to 37.5%. We have seen this with other books of business as well: a small percentage of accounts represent a large portion of the tail of the loss distribution.

	Portfolio Today	Optimized Portfolio
Number of accounts	173	157
Premium	$5,600	$5,200
Expenses	$1,700	$1,600
Expected Cat Loss	$500	$300
Expected Income	$3,400	$3,300
Loss at $99^{th}\% = F^{-1}(0.99)$	$28,600	$12,900
Capital Required	$23,200	$8,800
Return on Capital: ROC	14.7%	37.5%

Table 3: Portfolio before and after optimization. Unless otherwise noted, numbers are in (000).

Ideally, the portfolio manager should reprice these accounts upon renewal instead of terminating them. Although market conditions will determine the extent to which this is feasible, SmartWriter provides output on all the accounts targeted by the optimizer. Table 4 contains information for one of these accounts.

	Account A
Premium	$20
Expenses	$6
Expected Cat Loss	$12
Expected Profit	**$2**
Loss at $99^{th}\% = F^{-1}(0.99)$	$780
Capital Required	$740
Return on Capital: ROC	**0.3%**
Ret. on Marginal: ROMAC	**0.4%**
Premium needed to meet 15% ROC hurdle	**$150**
Premium needed to meet 15% ROMAC hurdle	**$145**

Table 4: Account targeted for removal or repricing by optimizer.

For this example, the premium needed to meet the stand alone return on capital hurdle of 15% is $150,000, much greater than the current premium of $20,000. Repricing is most likely not an option for this account, but for examples where the current ROC is closer to the hurdle rate, repricing can be viable.

5.1 Portfolio Expansion

As with the pruning portfolio example above, portfolio financials are available before and after optimization. Rather than repeat the above tables, we display the graphical output available from SmartWriter. Since the analysis was conducted at a zip code level, financials can be displayed in map form for quick understanding. We show an example below.

Figure 1 shows profitability by zip code, if each zip code is evaluated on a stand alone basis, for the San Francisco Bay area. Dark green indicates zip codes with a high expected ROC per home, light green less profitable, and red low profitability. These maps can be generated for expected income, marginal capital, and for the results of the optimization: optimal concentration by zip code. For confidentiality reasons, we do not give the recommended map for concentration, but it overlaps the map below to a large extent. Most zips that are low profitability the optimizer recommends moving away from, and for zips with high profitability, the optimizer recommends a greater penetration. The optimizer takes into account, however, the problems with overproducing in a number of closely located zip codes which all may be affected by the same earthquake.

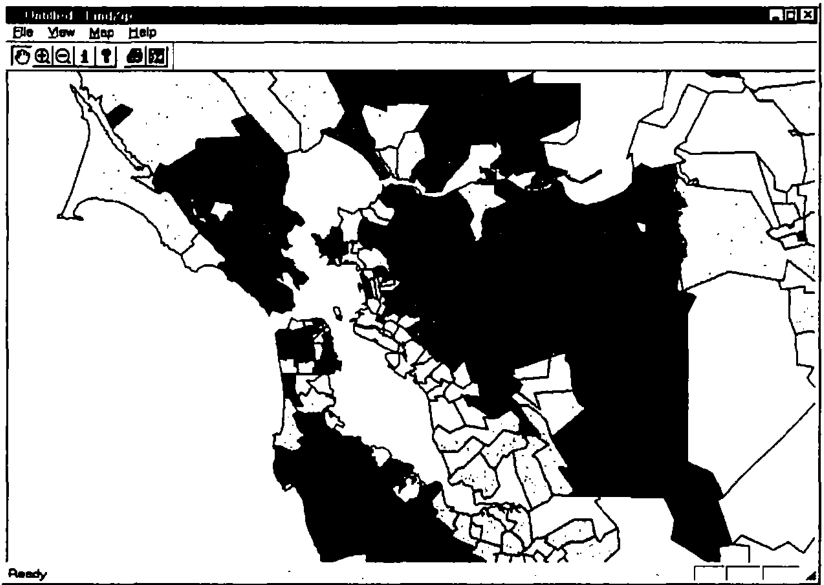

Figure 1: Expected return on capital by zip code for the San Francisco bay area. Dark green indicates most profitable zip codes, red indicates poor performing zip codes.

6. Next Steps

The portfolio management system can be readily extended to account for overlapping risks across business lines and asset investment categories. The concept is to develop a price of risk for each product-location under each scenario at each time period. These prices are available directly from the optimal decision variables for the strategic planning system. See Mulvey et al. (1998).

Ideally, one would like to link the liability decision with the asset investment strategy. In this paper we focused on the day-to-day underwriting decisions and take the asset return as a fixed input. In the future, one could tailor the asset portfolio in conjunction with the liability portfolio, such as purchasing catastrophe options or catastrophe-linked bonds for the property business line.

Another extension is the addition of multi-year contracts. As the catastrophe market continues to soften, these contracts may become more desirable for insurer and insured: They provide price protection for both parties. These can be linked with capital market projections which produce a range of possibilities (scenarios) a number of years ahead, such as the Towers Perrin CAP:Link system.

Finally, reinsurance decisions can be directly integrated into the optimization model. A desired profit distribution could be entered along with the current portfolio and a range of reinsurance options and treaties, and the optimizer would choose the best reinsurance options to match the desired profit distribution as closely as possible.

References

Bell, D., 1995, "Risk, Return, and Utility", *Management Science*, **41**, 23-30.

Berger, A., 1995, "Large Scale Stochastic Optimization with Applications to Finance," Ph.D. Thesis, Princeton University.

Glover, F., 1989, "Tabu Search – Part I", *ORSA Journal on Computing*, 1, 190-206.

Glover, F., 1990, "Tabu Search – a Tutorial", *Interfaces*, **20**, 74-94.

Glover, F., J. M. Mulvey, and K. Hoyland, 1996, "Solving Dynamic Stochastic Control Problems in Finance Using Tabu Search with Variable Scaling", in Meta-Heuristics: Theory and Applications, I. H. Osman and J. P. Kelly (eds.), Boston (MA): Kluwer Academic Publishers.

Mulvey, J., S. Correnti, and J. Lummis, 1997, "Total Integrated Risk Management: Insurance Elements," Princeton University Technical Report SOR-97-2.

Mulvey, J., C. Madsen, and F. Morin, 1998, "Linking Strategic and Tactical Planning Systems," DFA – Applications and Uses.

Mulvey, J., D. Rosenbaum, and B. Shetty, 1996, "Integrated Parameter Estimation in Stochastic Scenario Generation Systems", Princeton University Technical Report: SOR-96-15.