Optimal insurance risk allocation with steepest ascent and genetic algorithms

SiewMun Ha
Decision Analytics, AIR Worldwide, Boston, Massachusetts, USA

Received 23 October 2012
Revised 1 November 2012
5 November 2012
Accepted 12 November 2012

Abstract

Purpose – Enhanced risk management through the application of mathematical optimization is the next competitive-advantage frontier for the primary-insurance industry. The widespread adoption of catastrophe models for risk management provides the opportunity to exploit mathematical optimization techniques to achieve superior financial results over traditional methods of risk allocation. The purpose of this paper is to conduct a numerical experiment to evaluate the relative performances of the steepest-ascent method and genetic algorithm in the solution of an optimal risk-allocation problem in primary-insurance portfolio management.

Design/methodology/approach – The performance of two well-established optimization methods – steepest ascent and genetic algorithm – are evaluated by applying them to solve the problem of minimizing the catastrophe risk of a US book of policies while concurrently maintaining a minimum level of return.

Findings – The steepest-ascent method was found to be functionally dependent on, but not overly sensitive to, choice of initial starting policy. The genetic algorithm produced a superior solution to the steepest-ascent method at the cost of increased computation time.

Originality/value – The results provide practical guidelines for algorithm selection and implementation for the reader interested in constructing an optimal insurance portfolio from a set of available policies.

Keywords Insurance, Genetic algorithms, Risk analysis, Portfolio optimization, Catastrophe risk allocation, Steepest ascent

Paper type Research paper

Introduction

Primary-insurance companies routinely provide catastrophe insurance coverage for properties located in geographic areas subjected to low-frequency but high-severity natural perils such as hurricanes and earthquakes. Because the historical data for such perils are limited, these companies rely on computer simulations known as catastrophe (cat) models to provide scientifically-sound and statistically-valid data to help them quantify and manage the risks arising from the provision of such coverage. Cat-models augment the limited historical data with advanced scientific knowledge of the perils, and apply Monte-Carlo simulation to generate a statistically-sufficient number – typically 10,000 – of independent estimates of the annual losses arising from the occurrence of specified perils against given sets of exposures.

Aside from their primary utility in quantifying catastrophe risk for actuarial analyses, the cat-model outputs also encapsulate the detailed risk-correlation information that enables portfolio managers to optimally allocate the available risks to their portfolio so as to achieve specific financial objectives related to risk

The Journal of Risk Finance
Vol. 14 No. 2, 2013
pp. 129-139

1526-5943
DOI 10.1108/15265941311301170

minimization and/or return maximization. This follows from the fundamental principle established by Markowitz (1952) that the overall risk of a portfolio may be minimized by allocating individual risks that are uncorrelated with each other. Portfolio managers typically seek to achieve superior portfolio performance, and hence competitive advantage, through maximizing the return from their portfolio while concurrently minimizing their risk. Thus, a typical portfolio-optimization objective might be the maximization of the total portfolio return while limiting the attendant risk to a level consistent with the company's surplus. Alternatively, the objective may be defined as the minimization of the total portfolio risk while ensuring that the return meets a desired minimum threshold.

The task of selecting the optimal combination of risks to achieve a specified portfolio financial objective constitutes a combinatorial optimization problem belonging to the class of non-linear, multi-dimensional knapsack problems. The interested reader is referred to Martello and Toth (1990) and Kellerer *et al.* (2004) for additional information on knapsack problems. Zeng (2001) formulated the general insurance risk-management and portfolio-optimization framework, and presented a solution to the portfolio-optimization problem using Cauchy's (1847) Steepest-Ascent method. The Steepest-Ascent method operates by selecting an initial starting location in the solution space and then advancing to the nearest objective-function optimum by selecting intermediate steps that increase the objective-function value by the greatest extent at each step.

In Zeng's (2001) implementation of the Steepest-Ascent method, the optimal portfolio is constructed by sequentially adding policies based on their marginal impact on the objective function. The initial starting location is chosen, by design, as the policy with the best individual objective-function value. Hence, if the objective of the analysis is to minimize a given risk metric for the portfolio, then the initial starting point is the policy with the lowest individual risk metric.

When the solution space contains multiple local optima, it is intuitively obvious that, in such a method, the final solution is functionally dependent on the starting location, i.e. initial policy, that is chosen. This then raises the question as to whether the standard choice of the policy with the best individual objective-function value as the starting location produces the best final solution, or whether there exists another starting location which would produce a better final solution. For portfolios of moderate size, it is computationally feasible to evaluate this functional dependence, and determine the best solution, by exhaustively enumerating all the solutions relative to starting location, albeit at the cost of increased computation time. This in turn sets up the obvious solution-quality/computation-time tradeoff in solving this problem. Since the real-world portfolio management process typically contains a significant element of uncertainty, the pertinent issue for a portfolio manager is the practical significance of any improvement in solution quality relative to starting location, as opposed to the pure mathematical functional dependency. Hence, the first objective of this study is to evaluate the sensitivity and practical significance of the Steepest-Ascent method relative to its initial starting location by applying it to optimize a typical book of insurance policies.

An alternate method to solve the portfolio-optimization problem is the heuristic search Genetic Algorithm devised by Holland (1975), and described in detail by Whitley (1993) and Bäck (1996). The Genetic Algorithm represents a fundamentally different approach to solving the portfolio-optimization problem whereby it simulates the natural-selection process found in evolutionary biology to find a near-optimal solution.

In the Genetic Algorithm, a population of candidate solutions to the problem are coded as genetic strings and mated with each other to produce offspring which are then randomly mutated and subjected to natural selection and culling through a "predatory" objective function. Surviving solutions with superior fitness, i.e. objective-function values, are allowed to reproduce and repeat the process. The optimal solution emerges after a suitable number of generations as a result of the simulated evolutionary process. In contrast to the Steepest-Ascent method, the Genetic Algorithm is, in principle, able to overcome the problem of convergence to local optima by virtue of its stochastic design and operation. The current portfolio-optimization problem also lends itself easily to solution by the Genetic Algorithm as the decision vector is naturally coded as a genetic string of 1's and 0's and therefore requires no additional parameterization effort.

However, while the Genetic Algorithm offers the promise of finding a better solution, it is computationally more intensive than the Steepest-Ascent method and it thus potentially further extends the solution-quality/computation-time tradeoff. Furthermore, the stochastic nature of the Genetic Algorithm dictates that any single run of the algorithm effectively produces a solution that is a random drawing from a population of possible solutions, the parameters of which are dependent on the "tuning" optimality of the algorithm, that is the degree to which variables such as population size, parent/offspring ratio, mutation rate, crossover method and selection method have been optimized relative to the problem at hand. Hence, a second objective of this study is to examine the performance of the Genetic Algorithm relative to the Steepest-Ascent for the insurance portfolio-optimization problem. Performance in this case is measured by quality of final solution as well as computation time. Both methods were evaluated by applying them to find the optimal portfolio from a set of 500 insurance policies at multiple US locations subject to multiple perils.

Problem formulation

Consider a universe of N^* distinct insurance policies where each policy provides catastrophe coverage for one or more properties that may be located at a single location or distributed spatially. We wish to construct an optimal portfolio by judiciously selecting a subset of policies from this universe such that the overall portfolio risk, as measured by the 99 percent exceedance probability tail value at risk (TVaR), is minimized, while concurrently ensuring that the annual portfolio premium meets a specified minimum acceptable level of return. Assume that the 10,000-year cat-model loss vector and annual premium available for each policy are known.

We first define the following for the analysis:

$$L_{ij} = j\text{th annual loss estimate for policy i}; \quad \forall i = 1, 2, \ldots N^*; \; j = 1, 2, \ldots 10,000 \tag{1}$$

$$\mathbf{L}_i = \text{loss vector for policy i} = [L_{i1}, L_{i2}, \ldots L_{i10,000}]; \quad \forall i = 1, 2, \ldots N^* \tag{2}$$

$$P_i = \text{annual premium available for policy i}; \quad \forall i = 1, 2, \ldots N^* \tag{3}$$

$$X_i = \text{decision vector} = \begin{cases} 1, & \text{if policy i is allocated to the portfolio} \\ 0, & \text{if policy i is not allocated to the portfolio} \end{cases}; \quad \forall i = 1, 2, \ldots N^* \tag{4}$$

It is also convenient to define a sorting operator $\vec{\mathbf{S}}$ such that, for any real vector Z_k; $\forall k = 1,2, \ldots N$, $\vec{\mathbf{S}}\,[Z_k]$ produces a sorted vector $\mathbf{Z}_k^*$ with the property that: $\mathbf{Z}_1^* \geq \mathbf{Z}_2^* \geq \ldots \geq \mathbf{Z}_N^*$.

The total portfolio premium is then simply the sum of the premiums from the allocated policies:

$$\mathbf{P} = \sum_{i=1}^{N^*} (\mathbf{X_i P_i}) \tag{5}$$

The objective function is the portfolio TVaR at the 99 percent exceedance probability, defined by:

$$\mathbf{T} = \frac{1}{100} \sum_{j=1}^{100} \mathbf{L}_j^*, \tag{6}$$

where $\mathbf{L}_j^*$ is the *j*th element of the sorted total portfolio loss vector:

$$\mathbf{L}^* = \vec{\mathbf{S}} \left[\sum_{i=1}^{N^*} (\mathbf{X}_i \mathbf{L}_i) \right] \tag{7}$$

Consistent with Markowitz's (1952) results, this mathematical framework functionally links T to the degree of correlation between the individual risks through the $\sum_{i=1}^{N^*} (\mathbf{X}_i \mathbf{L}_i)$ term in equation (7) where the tail risk in $\mathbf{L}^*$ is reduced if uncorrelated individual risks $\mathbf{L_i}$ are selected for the portfolio.

Finally, the parameter that determines the initial starting policy for the Steepest-Ascent method is the individual policy TVaR, defined for the *i*th policy as:

$$\mathbf{T}_i = \frac{1}{100} \sum_{j=1}^{100} \mathbf{L}_{ij}^*, \tag{8}$$

where $\mathbf{L}_{\mathbf{ij}}^*$ is the *j*th element of the sorted *i*th policy loss vector:

$$\mathbf{L_i^*} = \vec{\mathbf{S}}[\mathbf{L_i}] \tag{9}$$

The fundamental portfolio-optimization problem is then formally defined as:

$$\text{Find } X_i \text{ to minimize } \mathbf{T}(\mathbf{X}_i)$$

$$\text{Subject to : } \mathbf{P}(\mathbf{X}_i) \geq \mathbf{C}$$

Experimental procedure

The data used in the current study consisted of a universe of $N^* = 500$ insurance policies located at multiple locations throughout the USA and subject to multiple catastrophe perils including earthquakes, hurricanes, severe thunderstorms, and winter storms. The total premium from the universe is $P^* = \$26.07M$ and the TVaR at the 99 percent exceedance probability is $T^* = \$104.2M$. For the current optimization exercise, the minimum acceptable portfolio premium was set at C = \$8.5M, approximately a third of the total premium from the universe.

The optimization problem was first solved by the standard Steepest-Ascent method described by Zeng (2001) which proceeds by sequentially adding policies to the portfolio based on the marginal impact of each policy on the objective function. This standard Steepest-Ascent method is denoted here as SA-1. Per this method, the first policy selected for the portfolio, corresponding to the initial starting location in the solution space, was the one with the minimum individual TVaR T_1. However, as previously discussed, the quality of the final solution is potentially sensitive to this choice of the first policy, especially if the solution space contains multiple local optima. This putative sensitivity was tested by exhaustively enumerating all possible solutions with respect to the initial policy choice. That is, each of the 500 policies in the universe was selected in turn as the initial policy, thereby producing a population of 500 distinct solutions corresponding to the 500 different initial starting policies. The best solution found among the 500 runs was then selected as the final solution. This exhaustive enumeration variant of SA-1 is denoted as SA-2.

The Genetic Algorithm described by Whitley (1993) and Bäck (1996) was implemented with the parameters, determined through trial and error, detailed in Table I. To account for its stochastic nature, the algorithm was run with $N_G = 200$ replications to obtain a reasonable sample size to evaluate its performance. The average result over the 200 runs, denoted as GA-1, is the expected value, i.e. an unbiased statistical estimate, for a "typical" single run of the Genetic Algorithm. Performance wise, we may thus regard GA-1 as representative of the result that would be obtained if we ran the algorithm only once. The best result over the 200 runs, denoted as GA-2, was also examined to gauge the run-to-run variation of the Genetic Algorithm. Table I summarizes the key characteristics of the four methods SA-1, SA-2, GA-1 and GA-2.

All the solution methods were programmed in MATLAB and executed on a Dell Latitude notebook computer with a 2.2 GHz Intel Core i7-2720QM processor and 8 GB of memory.

Experimental results

The performance metrics for all the solution methods are summarized in Table II and Figure 1. One measure of the effectiveness of each method is the improvement in the

Algorithm	Description
Steepest-Ascent 1 (SA-1)	Zeng (2001). Start with the policy with the lowest TVaR. Sequentially add policies based on the smallest marginal TVaR until the premium constraint is satisfied
Steepest-Ascent 2 (SA-2)	Variation of SA-1 with exhaustive enumeration of solutions with respect to the starting policy. Select the best final solution from the 500 solutions generated
Genetic Algorithm 1 (GA-1)	Whitley (1993) and Bäck (1996). 100 parents, 700 offspring, single-point crossover, mutation rate = 1 percent, elitist selection, stopping criterion: 250 iterations without improvement. Average solution from 200 runs
Genetic Algorithm2 (GA-2)	Whitley (1993) and Bäck (1996). 100 parents, 700 offspring, single-point crossover, mutation rate = 1 percent, elitist selection, stopping criterion: 250 iterations without improvement. Best solution from 200 runs

Table I. Steepest-Ascent and Genetic-Algorithm summary

TVaR/premium ratio (T/P) for the selected portfolio versus the corresponding ratio for the underlying universe of policies. By allocating policies selectively to improve the TVaR objective function, all the methods are able to deliver portfolios with superior TVaR/premium ratios compared to the universe, with T/P decreasing monotonically from 4.00 for the universe to 1.38, 1.22, 0.804 and 0.779 for SA-1, SA-2, GA1 and GA-2, respectively. On a practical level, we are interested in determining if the portfolio TVaR obtained by the Steepest-Ascent method is functionally dependent on the initial starting policy chosen, and the degree of sensitivity of that functional dependence. To answer these questions, we examine the population of solutions produced by SA-2, shown in Figure 2. A visual inspection reveals that a range of solutions is produced by the different the initial starting policies, and hence confirms the initial premise that there is a functional dependence between the portfolio TVaR obtained and the initial starting policy. However, the sensitivity is quite weak. The population has a coefficient of variation of only 2.0 percent and the entire population is located within a ±15 percent range of the mean. There is also a strong clustering around the mean with

Table II.
Performance metrics, all methods

Metric	Steepest-Ascent 1	Steepest-Ascent 2	Genetic Algorithm 1	Genetic Algorithm 2	Universe
Portfolio TVaR (T/T^*)	0.113	0.100	0.066	0.064	1.00
Portfolio premium (P/P^*)	0.327	0.327	0.326	0.326	1.00
TVaR/premium (T/P)	1.38	1.22	0.804	0.779	4.00
Policy count (N/N^*)	0.588	0.560	0.370	0.384	1.00
Computation time (t/t^*)	1.00	499	3.92	785	–

Notes: T^* = \$104.2M; P^* = \$26.1M; N^* = 500; t^* = 283s

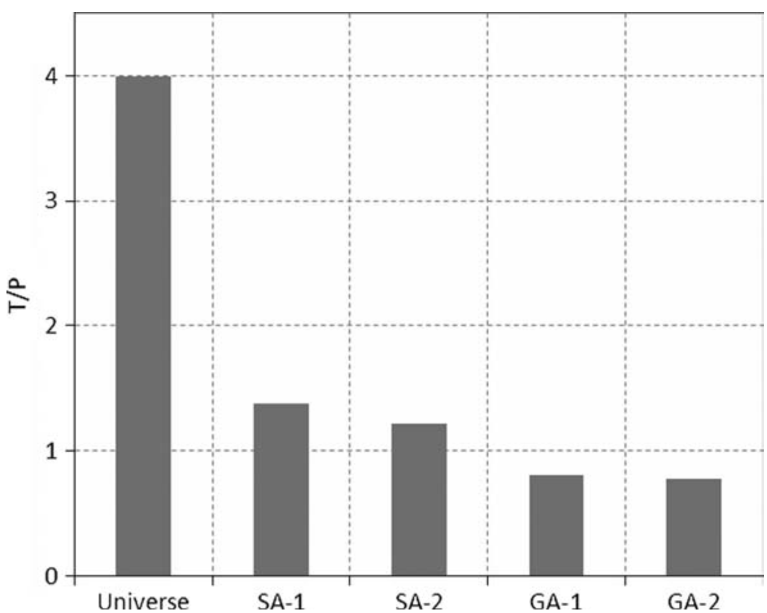

Figure 1.
TVaR/premium ratios, all methods

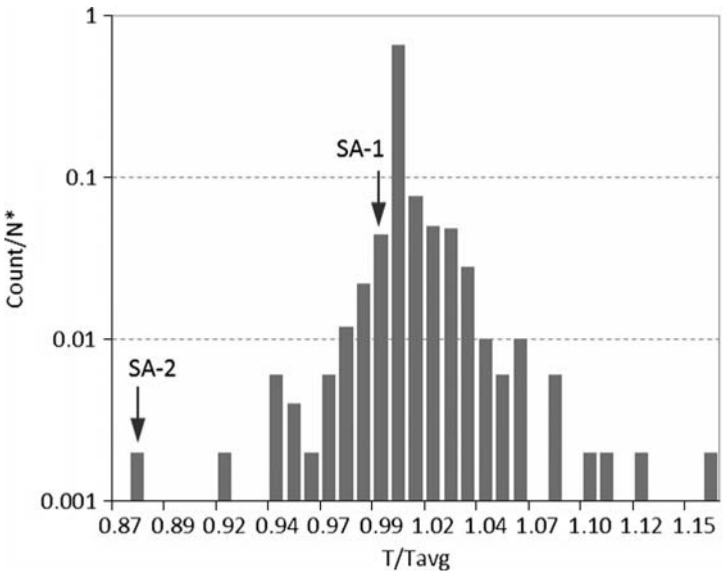

Figure 2.
SA-2 TVaR distribution

65.8 percent of the solutions are located within a range of mean ± 1 percent. Both SA-1 and SA-2 deliver premiums that are just above the $8.5M constraint, but SA-2 produces a portfolio TVaR that is 12.0 percent lower than SA-1.

While there is a distinct improvement in TVaR between SA-1 and SA-2, the portfolio composition itself changes very little between the two methods. The portfolio compositions are displayed graphically as grids in Figure 3(a) and (b), where a grey element indicates a policy that has been allocated to the portfolio, and a white element otherwise. The close similarity of the two portfolio compositions is easily identifiable from a visual inspection of the two grids. 485 of the 500 policies in the universe (97 percent) are allocated identically between the two portfolios and only 15 policies (3 percent) are allocated differently. These are highlighted in Figure 3(a) and (b). Of this set of 15 non-common policies, SA-1 has 14 that are not allocated to SA-2, and SA-2 has one that is not allocated to SA-1.

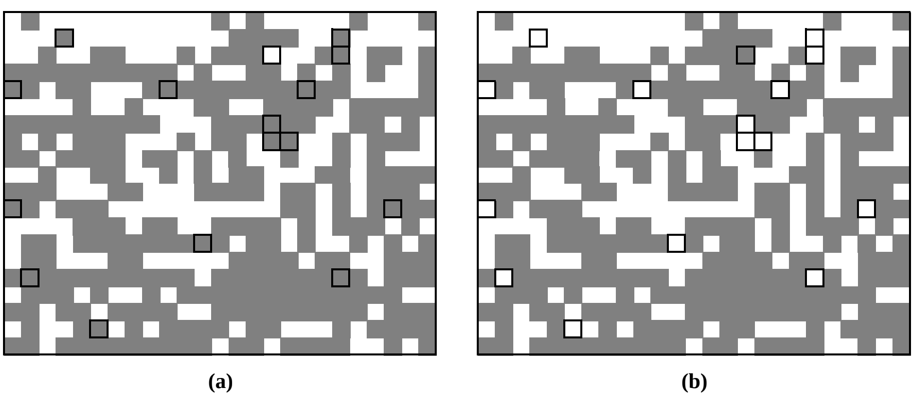

Figure 3.
(a) SA-1 portfolio composition; (b) SA-2 portfolio composition

Additional insight into the working of the Steepest-Ascent method may be gained by examining the relationship between the final portfolio TVaR, T, and the initial starting-policy TVaR, T_1. Figure 4 shows the correlation between T and T_1 over the 500 solutions from SA-2, plotted in increasing order of T_1. The variation in the portfolio TVaR exhibits a strong correlation with the magnitude of the starting-policy TVaR. SA-1, the standard application of the Steepest-Ascent method, begins with the policy with the lowest TVaR, corresponding to the first point on the left of the chart. A visual inspection indicates that the standard method produces a portfolio TVaR that is approximately average with respect to the 500 possible solutions. It therefore represents the "safe" option in terms of avoiding the extreme solutions, good and bad, in the population. On the other hand, counter intuitively selecting an initial policy with a large individual TVaR has the potential to produce a final portfolio TVaR that is much higher or lower than average.

Finally, we note that while SA-2 gives a superior TVaR result relative to SA-1, its strategy of exhaustive enumeration of all possible solutions with respect to starting policy forces an obvious tradeoff with respect to computation time. The exhaustive enumeration process dictates that the ratio of SA-2 computation time to that of SA-1 scales approximately with the size of the universe of policies (N^*) that is being analyzed. For the current problem, $N^* = 500$, and hence we have $t/t^* \approx 500$, where t and t^* are the computation times for SA-2 and SA-1, respectively. This linear increase with N^* imposes a practical limit on the scalability of applying SA-2 as a method for obtaining an improved solution relative to SA-1.

The solution-quality/computation-time tradeoff is further extended when the performances of the two instances of the Genetic Algorithm, GA-1 and GA-2, are compared against SA-1 and SA-2. The average Genetic-Algorithm solution GA-1 is significantly more effective than both SA-1 and SA-2 with a TVaR that is 42.0 and 34.2 percent lower, respectively, while offering an approximately equal premium of \$8.5M versus \$8.53M for SA-1 and SA-2. This improved performance was obtained with a computation time that is 3.9 × that of SA-1 but only 0.0079 × that of SA-2. SA-1 and GA-1 therefore form a segment of the Pareto front for TVaR and computation time.

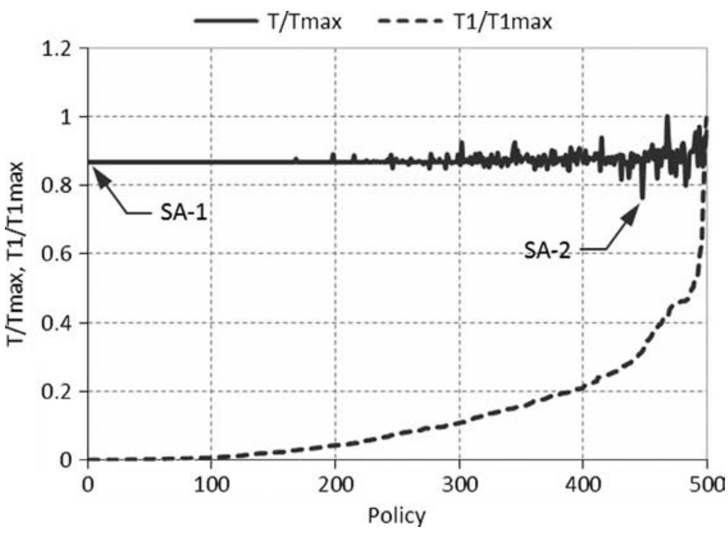

Figure 4.
Portfolio TVaR (T) and initial policy TVaR (T_1) correlation

However, SA-2, being inferior to GA-1 in both solution quality and computation time is not located on the front.

Examining the best solution from the 200 Genetic-Algorithm runs, we find that GA-2 provides only a marginal improvement over the average solution GA-1 at the cost of a greatly increased computation time. GA-2 TVaR is reduced by 3.16 percent relative to GA-1 with an attendant 200 × increase in computation time. GA-2 therefore extends the Pareto front for TVaR and computation time. The Pareto front for the TVaR/computation-time tradeoff for all the four methods examined in this study is shown in Figure 5. There is a large reduction in TVaR (42.0 percent) for a modest increase in computation time (3.9 ×) in proceeding from SA-1 to GA-1. On the other hand, there is only a modest reduction in TVaR (3.16 percent) for a large increase in computation time (200 ×) in proceeding from GA-1 to GA-2. Hence, the average Genetic-Algorithm solution GA-1, representing a typical "single-run" solution, is the best compromise between solution quality and computation time.

The overall variation in the run-to-run population in the Genetic Algorithm is small with a coefficient of variation of only 1.9 percent. Figure 6 shows the distribution of the portfolio TVaR, T, obtained over the 200 GA-2 runs. The population is located within the range of 98-106 percent of the mean and exhibits an asymmetric bi-modal distribution. This bi-modal structure was confirmed with an additional experiment with 1,000 runs on the same set of data. The run-to-run variation in this case is a function of the specific characteristics of the current problem as well as the effectiveness of the parameter tuning of the algorithm. A possible explanation for this bi-modal structure is that the solution space may contain two dominant local optima that trap the algorithm when the applied mutation rate is insufficiently aggressive to extend the solution search to a new region beyond them. The confirmation of this hypothesis is unfortunately beyond the scope of the current study.

Aside from the lower portfolio TVaR obtained, another notable difference between the Genetic-Algorithm solutions and their Steepest-Ascent counterparts is that

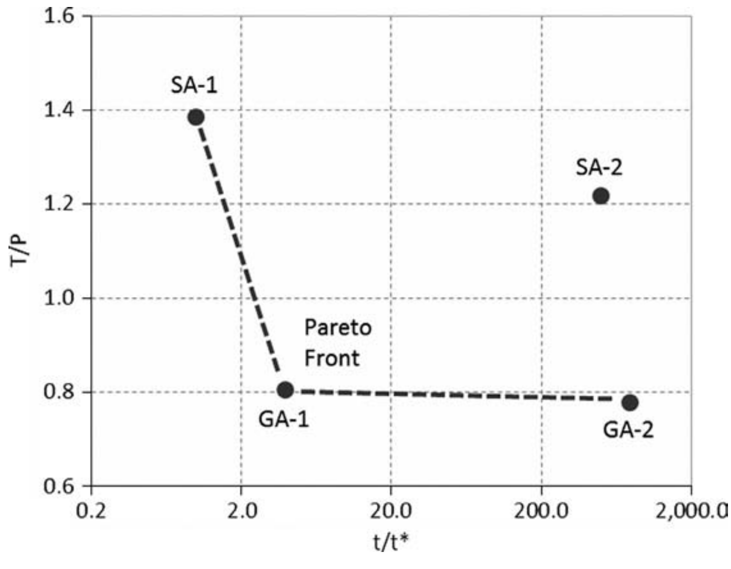

Figure 5.
TVaR/computation-time Pareto-optimal front

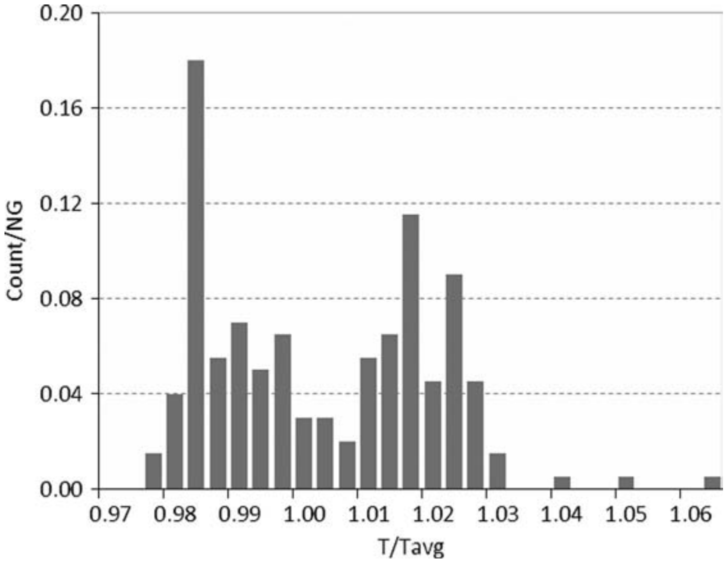

Figure 6.
GA-2 TVaR distribution

Note: $N_G = 200$

GA-1/GA-2 (185/192) have a distinctly smaller number of policies in their portfolios than SA-1/SA-2 (294/280). This is advantageous from a practical business perspective as, all else being equal, smaller portfolios are simpler to manage and have correspondingly lower overhead expenses such as underwriting and claims-processing costs.

Conclusions

We conducted a numerical experiment to evaluate the relative performances of the Steepest-Ascent method and Genetic Algorithm in the solution of an optimal risk-allocation problem in primary-insurance portfolio management.

The portfolio TVaR obtained by the Steepest-Ascent method is functionally dependent on the initial starting policy that is chosen. For the current problem representing a typical US insurance book of business, the sensitivity of this functional dependency is fairly low with a coefficient of variation of only 2 percent. The population of solutions is tightly distributed around the mean with the entire population located within a ± 15 percent range and 65.8 percent located within a ± 1 percent range.

The standard Steepest-Ascent method (Zeng, 2001) gives a TVaR that is approximately equal to the mean solution over all possible starting locations. The TVaR from the best solution over all possible stating locations is 12.0 percent lower than the standard solution. This is obtained at the cost of a $500\times$ increase in computation time, proportional to the size of the policy universe.

A "typical" average single run of the Genetic Algorithm (GA-1) achieves a significantly superior TVaR result to both SA-1 and SA-2. GA-1 TVaR is 42.0 and 34.2 percent lower relative to SA-1 and SA-2, respectively. This improved solution quality is achieved with a modest increase in computation time relative to SA-1 ($3.9\times$) and a much reduced computation time relative to SA-2 ($0.008\times$).

The run-to-run variation in the Genetic Algorithm, with the current parameter tuning, is small with a coefficient of variation of 1.9 percent for the population of solutions produced over 200 runs of the Genetic Algorithm. The population exhibits a bi-modal distribution, which indicates a potential interaction between the algorithm and local optima in the solution space. Further research is required to confirm this.

The Pareto-optimal front for solution-quality/computation-time tradeoff consists of SA-1, GA-1 and GA-2 but not SA-2. GA-1 offers the best practical balance between solution quality and computation time among the methods tested.

Further research is required to isolate the cause of the bi-modal distribution identified in the GA-2 results, improve optimization performance and extend the range of applicability of the current results. The literature suggests that higher-quality solutions are potentially achievable by a hybrid approach (Hacker *et al.*, 2002), whereby the Genetic Algorithm is utilized to explore the solution space and the Steepest-Ascent method is applied to locate the final optimum solutions in promising regions found by the Genetic Algorithm. Operational Genetic-Algorithm parameters such as population size, parent/offspring ratio, mutation rate, crossover method and selection method were determined by trial and error in the current exercise and there is potential to achieve improvements in solution quality and computation time through the optimization of these parameters. Finally, it would be useful to repeat the current experiment for a range of different exposures and perils to check the robustness of the conclusions across different portfolio types.

References

Bäck, T. (1996), *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*, Oxford University Press, New York, NY.

Cauchy, A. (1847), "Méthode générale pour la résolution des systèmes d'équations simultanées", *Comptes Rendus Hebd. Séeances Acad. Sci.*, Vol. 25, pp. 536-8.

Hacker, K.A., Eddy, J. and Lewis, K.E. (2002), "Efficient global optimization using hybrid genetic algorithms", paper presented at 9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Atlanta, GA.

Holland, J.H. (1975), *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, MI.

Kellerer, H., Pferschy, U. and Pisinger, D. (2004), *Knapsack Problems*, Springer, Berlin.

Markowitz, H.M. (1952), "Portfolio selection", *The Journal of Finance*, Vol. 7 No. 1, pp. 77-91.

Martello, S. and Toth, P. (1990), *Knapsack Problems: Algorithms and Computer Implementations*, Wiley, Chichester.

Whitley, D. (1993), *A Genetic Algorithm Tutorial*, Computer Science Department, Colorado State University, Fort Collins, CO.

Zeng, L. (2001), "Using cat models for optimal risk allocation of P&C liability portfolios", *The Journal of Risk Finance*, Vol. 2 No. 2, pp. 29-35.

About the author

SiewMun Ha is a Senior Risk Consultant at Decision Analytics, AIR Worldwide, Boston, MA, USA.