MARCH 2012

AIRCURRENTS: PORTFOLIO OPTIMIZATION FOR REINSURERS

BY SIEWMUN HA, PH.D.
EDITED BY NAN MA

EDITOR'S NOTE: *A previous AIRCurrent explored portfolio optimization techniques for primary insurance companies. In this article, Dr. SiewMun Ha (Senior Risk Consultant) discusses advanced optimization solutions for reinsurance companies.*

THE ARTICLE: Discusses the challenges in finding optimum participation levels in portfolios of reinsurance contracts and explains AIR's innovative approach.

HIGHLIGHTS: AIR's stochastic search algorithms are able to overcome the disadvantages of more traditional methods such as policy ranking and hill climbing to provide a superior capability of locating the global optimum, while taking into account multiple sources of uncertainty and business constraints.

INTRODUCTION

A common, if not universal, business objective for reinsurers is the maximization of a revenue metric such as premium (P), or minimization of a risk metric such as Tail Value at Risk (TVaR), or optimization of some composite metric such as P/TVaR, while simultaneously satisfying multiple constraints imposed by capital, market and business requirements.

This optimization is achieved through judicious control of the participation levels of the company's portfolio of contracts. The selection of the appropriate contract-participation levels to optimize the desired combination of metrics constitutes an *optimization problem* encapsulating two difficulties—one mathematical and the other computational—which render its solution process a non-trivial challenge.

The mathematical difficulty arises from the fact that the standard risk metrics are typically nonlinear with respect to the contract participations as a result of *partial correlations* between the modeled contract losses. This means that simulated events that cause high losses to one contract within a portfolio do not necessarily cause high losses to other contracts. This can be the case for a number of reasons, including geographic distance and dissimilarity in vulnerability of the underlying exposure. See the text box below for a simple example.

WHAT IS PARTIAL CORRELATION?

Consider a portfolio consisting of two reinsurance contracts with ten-year losses, L_1 = [35.2, 39.4, 90.8, 2.1, 43.7, 94.1, 18.6, 98.7, 39.6, 45.5] and L_2 = [43.9, 57.0, 35.9, 26.7, 42.2, 1.0, 3.3, 62.3, 64.8, 11.3], and participations X_1 and X_2, respectively. The combined portfolio loss **L** is then a linear combination of the individual contract losses and their respective participations: $\mathbf{L} = X_1 * L_1 + X_2 * L_2$

Suppose that the pertinent risk metric is the 70% TVaR. If the individual contract losses were perfectly correlated, the portfolio TVaR ($T_{70\%}$) would be a linear function of the individual contract TVaRs and their respective participations, that is: $T_{70\%} = X_1 * T_1 + X_2 * T_2$ where $T_1 = 94.5$ and $T_2 = 61.4$ are the 70% TVaRs for contracts 1 and 2, respectively. Restricting $X_2 = 1 - X_1$ for simplicity, this produces the linear function of $T_{70\%}$ against X_1 shown by the blue line in Figure 1. However, because the losses are generally only *partially correlated* between contracts in actual portfolios, TVaR is nonlinear with respect to X_1 and X_2, and must be evaluated for each separate combination of X_1 and X_2. This is done by first computing the combined losses **L**, ranking them, and then computing the TVaR from the resulting ranked losses. This process produces the convex green line in Figure 1.

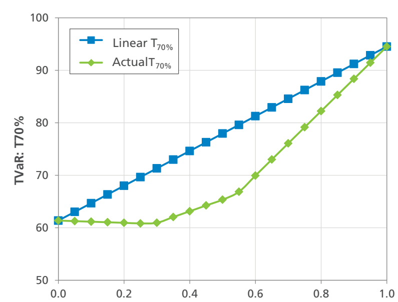

Figure 1. Portfolio 70% TVaR for two contracts with ten-year modeled losses

This nonlinearity greatly increases the mathematical complexity of the portfolio optimization analysis wherever TVaR or similar risk metrics are a component of either the objective function or the constraints.

The computational difficulty arises from the size of the space in which the search for the optimum solution must take place. In the above example, if we discretize the 0 – 100% potential participation range for each contract into 2-percentage-point steps, we obtain 100/2 = 50 discrete steps for each contract and a total of $50^2 = 2{,}500$ possible participation combinations for the portfolio. This is a trivially small solution space relative to today's computing power and it would be a simple matter to exhaustively compute the losses for every combination in order to identify the one with the minimum TVaR. However, the solution space grows exponentially with the number of contracts. A typical reinsurance portfolio can contain as many as 500 contracts, each with a 10,000-year loss vector, which produces a vastly larger solution space with $50^{500} = 3.1 \times 10^{849}$ possible combinations. This exceedingly large space, coupled with the greatly increased loss-vector length, makes it infeasible to exhaustively compute all the possible combinations. In fact, the required computation time with the fastest available computer would exceed the age of the universe by many orders of magnitude.

OPTIMIZATION

As a result of these mathematical and computational difficulties, reinsurance portfolio optimization demands a more systematic approach to efficiently locate the optimum in the solution space. Methods currently available range from simpler algorithms such as *policy ranking* and *hill climbing* to more advanced heuristic techniques such as *genetic algorithms* and *evolutionary search.*

In policy ranking, the contracts are ranked by the metric of interest and added to the portfolio one at a time, starting with the best, until the first constraint is reached. This is a simple and intuitive method of portfolio contract selection requiring only a minimal computational effort. However, it does not account for the less-than-perfect correlations between the modeled losses and will therefore most likely produce a suboptimal solution.

Hill climbing methods, such as *steepest ascent*, start at a given point in the solution space and explore the surrounding region to find locations that improve performance. It then moves to the best solution among those and repeats the process until it can no longer find any improvement in the local neighborhood. This method will generally deliver a superior result relative to policy ranking. However, its disadvantage is that the final solution critically depends on the initial starting point and it subsequently searches a progressively more restricted solution space in successive iterations until it converges to the nearest local optimum. Consequently, unless the global optimum fortuitously happens to also be the local optimum nearest to the selected starting point, a hill climbing analysis will also produce a suboptimal solution.

Stochastic search methods such as genetic algorithms and evolutionary search are the most advanced methods currently available and are particularly suited to problems of this nature. These algorithms were developed by mathematician John Holland in the 1970's to solve certain classes of mathematical optimization problems by simulating the natural-selection process in evolutionary biology (see Holland, 1992 for further reading). Genetic algorithms and evolutionary search operate on the following general framework:

1. **Initialization**: randomly generate an initial population of solutions
2. **Reproduction:** produce a new generation of solutions (offspring) from the initial population (parents) by randomly combining information from randomly selected pairs of parents
3. **Mutation:** induce random mutation in the offspring
4. **Ranking:** evaluate the fitness of each solution and rank them from best to worst
5. **Selection:** cull the least fit offspring and retain the most fit ones to serve as the parents for the next generation
6. **Iteration:** repeat Steps 2–5 until some exit criterion, such as maximum iteration count or fitness threshold level, is met
7. **Termination:** save the best solution found and exit

Stochastic search algorithms, by virtue of their stochastic nature, are able to overcome the disadvantages of the policy ranking and hill climbing methods to provide a superior capability of locating the global optimum. On the other hand, they are more computationally intensive than the simpler methods, and thus more time consuming. (A more detailed description of genetic algorithms may be found in Goldberg, 1989.)

AIR utilizes the entire range of available optimization methods to efficiently find the best solutions. The simpler methods complement the more advanced ones by providing a first-order approximation of the solution that can then be used to improve the starting point for the more advanced methods.

UNCERTAINTY ANALYSIS

The challenge of finding, attaining and maintaining an optimal reinsurance portfolio is further complicated by multiple sources of uncertainty in the process. These uncertainties affect both the location of the optimal solution itself and the ability of the portfolio manager to attain the optimal operating point. Consequently, it is not sufficient to merely identify the optimal solution; a comprehensive reinsurance portfolio management approach must also account for these uncertainties in the analysis so that the resulting variation around the optimal solution can be identified and quantified to provide the portfolio manager with a complete stochastic view of the process of adjusting the portfolio from its initial state to its optimal state.

Some examples Include:

- **Uncertainty in the choice of risk measure due to operational ambiguity**, which arises from multiple possible choices in the problem formulation. The sensitivity to risk measure choice may be evaluated by reanalyzing the portfolio with different risk measures and evaluating the change in the optimal solution with each selected risk measure.
- **Uncertainty in attainable contract participations due to market and business restrictions,** which represents a form of noise over the control of the portfolio and makes it improbable that the exact optimal solution can actually be attained. This naturally raises the question: *How close can I get to the optimal solution?* The sensitivity to this noise may be evaluated by conducting a Monte Carlo simulation to estimate the joint confidence intervals for the premium and losses around the optimal solution.
- **Uncertainty in modeled losses due to catastrophe model limitations**, which represents a form of noise over the inputs to the problem and it results in a corresponding uncertainty in the losses at the optimal solution. The sensitivity to this noise may again be evaluated through a Monte Carlo simulation to estimate the loss confidence interval at the optimal solution.

SOLUTION ROBUSTNESS

The solution space to an optimization problem may, in general, contain multiple local optima in addition to the global optimum (Figure 2). From a practical business perspective, the absolute global optimum may not necessarily be the best solution if prevalent uncertainties in participations and modeled losses make it difficult to attain or maintain the exact desired optimum operating point. Such a solution is not *robust*, meaning that a small deviation from the optimum operating participations can result in a significant performance degradation in the desired metric (Point A in Figure 2).

It may thus be better to accept a lower level of performance (Point B in Figure 2) in exchange for an increased level of robustness. To make an informed decision on the performance-robustness tradeoff, it is necessary to map the solution space to identify the location of each optimum and its surrounding topography. This may be accomplished by employing a *Tabu search* algorithm in conjunction with one of the standard optimization algorithms discussed earlier. The Tabu search process works by systematically directing the coupled optimization algorithm to search the solution space for multiple optima over successive iterations to produce the desired map. Hillier and Lieberman (2010) provide additional information on Tabu search.

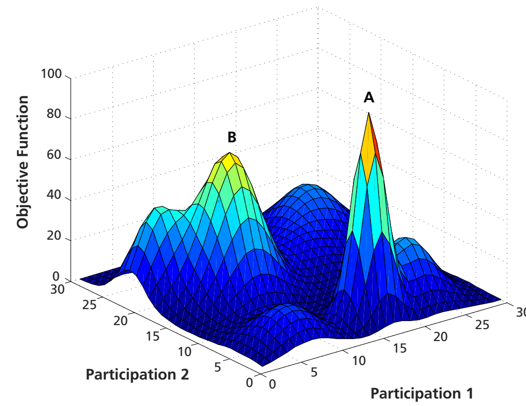

Figure 2. Solution robustness: global optimum (A) vs. local optimum (B)

CASE STUDY

The optimization and uncertainty concepts discussed in the previous sections are illustrated here in a case study of a hypothetical reinsurer with a small portfolio of U.S. contracts.

INTRODUCTION

Hutan Re, a major reinsurance company, possesses a small U.S. reinsurance portfolio with 12 regional contracts covering hurricane, severe thunderstorm and winter storm risks. The basic parameters of the current portfolio are summarized in Table 1. Current contract participations range from 4.1% to 9.7% and business/market restrictions impose upper and lower limits on individual contract participations. Current premiums for the contracts range from $0.5M to $37M. Participations for Contracts 3, 7 and 8, which belong to clients who are considered critical, cannot be altered. *Hutan* uses tail value at risk (TVaR) at the 1% exceedance probability as its primary risk metric. The current portfolio's total premium and TVaR are $101.5M and $49.4M, respectively.

Table 1. *Hutan Re* baseline reinsurance portfolio

Contract ID	Current Participation	Minimum Participation	Maximum Participation	Premium ($M)
1	4.07%	1.9%	9.1%	3.0
2	6.19%	1.2%	11.4%	4.7
3	8.06%	8.1%	8.1%	3.3
4	5.41%	1.2%	9.4%	20.9
5	6.59%	4.7%	10.4%	5.0
6	9.02%	4.0%	11.9%	5.9
7	4.84%	4.8%	4.8%	0.59
8	7.07%	7.1%	7.1%	0.51
9	8.56%	2.1%	10.9%	37.0
10	9.65%	4.0%	11.5%	11.2
11	6.63%	3.1%	10.6%	1.1
12	4.74%	2.7%	8.0%	8.2

Hutan is about to renew its portfolio and the portfolio manager is interested in investigating the following issues to inform his decisions:

1. How much additional premium can *Hutan* obtain if the current portfolio is optimized with the risk capped at its current level?
2. Where, in the TVaR-premium space, does the best risk-return tradeoff occur?
3. What if, contrary to the current paradigm, the critical clients are assumed to be NOT critical and the participation constraints on Contracts 3, 7 and 8 are relaxed?
4. How does the optimum portfolio change if the modeled losses for contracts 1, 4, 8 and 11 increase by 30% due to catastrophe model updates?
5. What if ALL the modeled losses contain an inherent uncertainty in the range of ±30%?
6. What is the effect of a participation uncertainty of ±2 percentage points (PP) on the optimal solution?

PROBLEM FORMULATION

The issues posed by scenarios 1 – 4 constitutes an optimization problem to maximize the premium (objective function) by controlling the contract participations (decision variables), while subject to constraints on the participations themselves and also the TVaR, which is a function of the participations and modeled losses. The problem may be stated formally as:

Maximize:

	$P_T = \Sigma(X_i * P_i)$	Objective function
Subject to		
	$a_i < X_i < b_i$	Participation constraints
	$TVaR(L) < T_{max}$	TVaR constraint
	$L = \Sigma(X_i * L_i)$	Portfolio loss function

Where

P_T = total portfolio premium	TVaR(**L**) = portfolio TVaR function
X_i = contract i participation level	T_{max} = maximum allowable TVaR
P_i = contract i maximum premium	**L** = portfolio loss vector
a_i = contract i minimum participation	$\mathbf{L}_i$ = contract i loss vector
b_i = contract i maximum participation	

ANALYSIS RESULTS

AIR programmed a customized evolutionary search algorithm to answer the questions posed by *Hutan's* portfolio manager. The uncertainty analyses for scenarios 5 – 6 were addressed through a Monte Carlo simulation program coupled with the main optimization algorithm. The results for the first three considerations are plotted on a premium-TVaR plane in Figure 3. The current baseline portfolio is marked by the red "X". Optimizing the baseline portfolio with the risk capped at its original level ($49.4M) produces an optimized portfolio with an additional ΔP = $26.7M in premium.

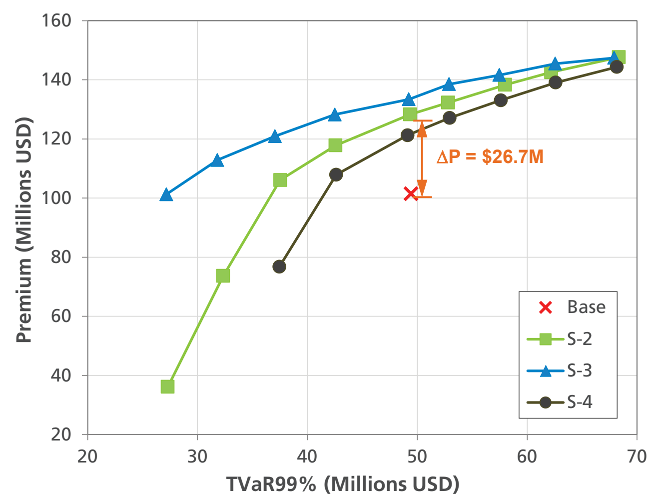

Figure 3. Reinsurance portfolio optimization scenario results

The baseline Pareto-optimal front is obtained by repeating the optimization over the feasible range of risk levels and is represented by the green line (S-2). This represents the range of the best possible risk-premium tradeoffs for the baseline portfolio. It is pertinent to note that the concave curvature of the front implies that the best risk-premium tradeoff occurs at the lower end of the risk range and it becomes progressively worse as more risk is written.

Relaxing the critical-client assumption on Contracts 3, 7 and 8 shifts the Pareto-optimal front upwards to the position of the blue line (S-3) in Figure 3. This indicates that more premiums are obtainable at every level of risk by exploiting the extra degrees of participation freedom made available by the relaxation of the constraints. At the baseline TVaR of $49.4M, the improved front (S-3) offers an increased premium of $5.1M compared to the optimized baseline (S-2). This represents the effective opportunity cost of maintaining the status of these clients as critical.

For scenario 4, repeating the optimization using +30% revised loss estimates for Contracts 1, 4, 8 and 11 shifts the Pareto-optimal front down to the position of the black line (S-4) in Figure 3. In this case, the increased losses force a reduction in the written risk to satisfy the TVaR constraint and consequently reduces the obtainable premium along the entire range of the Pareto-optimal front. At the baseline TVaR of $49.4M, this reduction amounts to $7.0M, representing the opportunity cost imposed by the modeled loss uncertainty.

For scenario 5, a Monte Carlo simulation was applied to the computation of the TVaR along the baseline Pareto-optimal front to evaluate the effect of a random ± 30% uncertainty in the complete set of modeled losses with contract participations kept constant. The modeled loss uncertainty produces a corresponding distribution of TVaR values at each point on the Pareto optimal front. Figure 4 shows one such distribution around the TVaR = $49.4M point, corresponding to the optimized baseline portfolio. Note the non-intuitive result that the distribution is *asymmetrical* around the baseline TVaR value, with the baseline TVaR itself located near the lower tail end of the distribution. This stems from the mathematical structure of the TVaR definition, which incorporates an inherent positive bias with respect to *uncorrelated random variation* in the modeled losses.

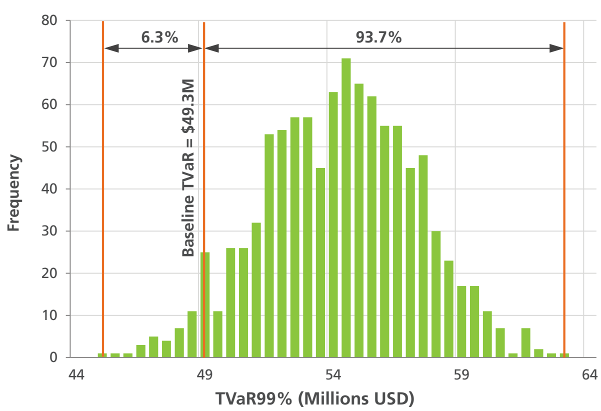

Figure 4. Baseline optimal solution: ± 30% modeled-loss uncertainty TVaR confidence interval

In this particular case, there is a 93.7% probability that the modeled loss uncertainty will produce a TVaR that is up to $13.4M higher than the baseline value. The maximum probability of the distribution, i.e., most probable TVaR value, corresponds to TVaR ≈ $54.5M, an increase of $5.1M over the baseline value. Repeating the Monte Carlo simulation along the entire Pareto-optimal front produces the corresponding confidence interval shown in Figure 5.

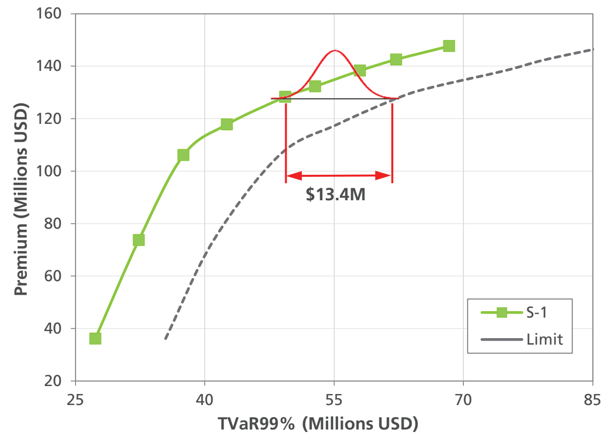

Figure 5. Baseline pareto-optimal front: ±30% modeled loss uncertainty TVaR confidence interval

Finally, for scenario 6, a second Monte Carlo simulation was applied to evaluate the effect of a ±2 percentage-point participation uncertainty on both TVaR and premium for the optimized baseline portfolio. Here, the participation uncertainty produces a corresponding two-dimensional joint probability distribution for TVaR and premium. This is illustrated in Figure 6 with the color-coded confidence intervals around the optimized baseline portfolio marked by the blue "X". Figure 7 shows the three-dimensional perspective view of the same confidence intervals. Each zone's color code correlates with the probability of the portfolio's final location in the premium-TVaR plane occurring in that zone as a result of participation uncertainty.

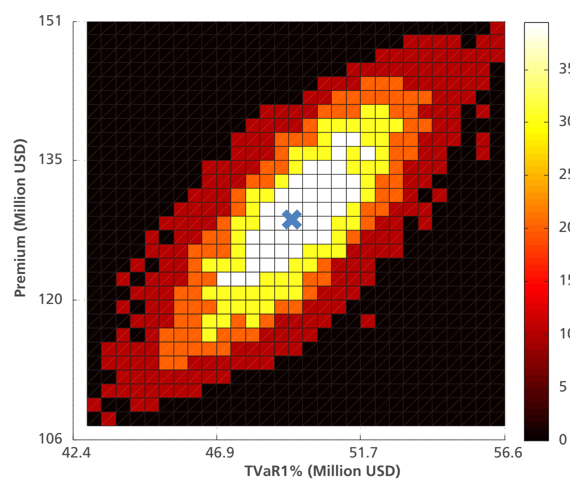

Figure 6. Baseline optimal solution: ±2 percentage-point participation uncertainty confidence interval

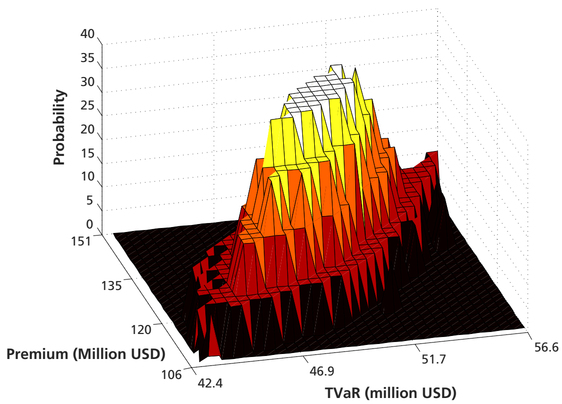

Figure 7. Perspective view: ±2 percentage-point participation uncertainty confidence interval

The individual probabilities for the white, yellow, orange, red and black zones are 40%, 30%, 20%, 10% and 0%, respectively. Mapping the zone probabilities to their respective TVaR and premium ranges then defines the following cumulative confidence intervals due to the participation uncertainty, as shown in Table 2.

Table 2. Cumulative confidence intervals due to participation uncertainty

Figure 6 Zone Color Code	TVaR Confidence Interval	Premium Confidence Interval	Confidence Level
White	$46.9M < T < $51.7M	$122M < P < $137M	40%
White + Yellow	$46.4M < T < $52.7M	$115M < P < $140M	70%
White + Yellow + Orange	$44.9M < T < $53.7M	$112M < P < $143M	90%
White + Yellow + Orange + Red	$42.5M < T < $56.6M	$106M < P < $151M	100%

CONCLUSION

Hutan Re gained the following insights from AIR's analysis of their portfolio:

1. Optimizing their current baseline portfolio delivers an additional $26.7M in premium for the same level of risk
2. The Pareto-optimal front, representing the range of best possible risk-return tradeoffs, occupies the range $27.3M < TVaR < $68.3M and $36.2M < P < $148M. Incremental risk delivers decreasing premium as more risk is written
3. Maintaining the critical clients at their current contract participation levels incurs an opportunity cost of $5.1M
4. A +30% revision in the modeled losses for Contracts 1, 4, 8 and 11 reduces the premium along the Pareto-optimal front, with the reduction increasing as less risk is written. The premium reduction amounts to $7.0M for the optimized baseline portfolio
5. A ±30% uncertainty in all the modeled losses produces a 93.7% asymmetrical confidence interval of $13.4M on the positive side of the optimized baseline TVaR value
6. A ±2 percentage-point uncertainty in the contract participations produces a two-dimensional confidence interval around the optimized baseline. At the 90% confidence level, $44.9M < TVaR < $53.7M and $112M < P < $143M

With the optimization framework and model established, *Hutan* is now also positioned to evaluate revised scenarios and/or combinations of scenarios as the need arises.

FURTHER READING

Holland J.H., *Adaptation in Natural and Artificial Systems,* A Bradford Book, 1992

Goldberg D.E., *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, 1989

Hillier F.S. & Lieberman G.J., *Introduction to Operations Research*, McGraw Hill, 2010

ABOUT AIR WORLDWIDE

AIR Worldwide (AIR) is the scientific leader and most respected provider of risk modeling software and consulting services. AIR founded the catastrophe modeling industry in 1987 and today models the risk from natural catastrophes and terrorism in more than 90 countries. More than 400 insurance, reinsurance, financial, corporate, and government clients rely on AIR software and services for catastrophe risk management, insurance-linked securities, detailed site-specific wind and seismic engineering analyses, agricultural risk management, and property replacement-cost valuation. AIR is a member of the Verisk Insurance Solutions group at Verisk Analytics (Nasdaq:VRSK) and is headquartered in Boston with additional offices in North America, Europe, and Asia. For more information, please visit **www. air-worldwide.com**.