Proceedings of the 2007 INFORMS Simulation Society Research Workshop

TWO-LEVEL SIMULATIONS FOR RISK MANAGEMENT

Hai Lan
Barry L. Nelson
Jeremy Staum

Dept. of Industrial Engineering & Management Sciences
Northwestern University
Evanston, IL 60208-3119, U.S.A.

ABSTRACT

Risk measurement involves estimating some functional of a loss distribution. This calls for nested simulation, in which risk factors are sampled at an outer level of simulation, while the inner level of simulation provides estimates of loss given each realization of the risk factors. We present a general method for providing a confidence interval for the risk measurement. It involves combining a confidence region for the losses with the confidence interval that would be used if the losses were known. This method could be very time-consuming, so we discuss ways of improving its efficiency. One is to choose the structure of the confidence region for the losses in a way that is tailored to the particular risk measure. Another involves varying the simulation effort expended on estimating the losses associated with different risk factors, even screening out many of the risk factors after slight simulation effort has revealed them to be unlikely to have much impact on the estimated risk measure. We will discuss two-level simulation in general and a specific procedure for estimating a confidence interval for tail conditional expectation.

1 INTRODUCTION

Every night, institutions trading in derivative securities measure the risk of their portfolios of derivatives. The Basel II regulatory framework (Bank of International Settlements 2004) encourages the calculation of capital requirements for banks on the basis of banks' internal risk assessments. These capital requirements against possible losses are a crucial aspect of regulatory efforts to prevent a cascade of defaults that could paralyze the global financial system. Risk measurement also allows firms to impose risk management policies that prevent limits on risk from being breached, helps to guide efforts to mitigate risk by hedging, and enables a policy of charging individual business units for risk capital, thus providing incentives for risks to be taken only when the expected rewards are sufficiently large.

A number of difficulties attend the practice of nightly firmwide risk management. One is data aggregation: a large financial institution may have positions in thousands of derivative securities, involving hundreds of underlying variables. Another is the difficulty of modeling future risks on the basis of presently available information. A third difficulty is purely computational and can be addressed by improved simulation methodology.

Risk measurement involves computing something about the distribution of possible losses on the portfolio over some time horizon. For example, the one-day 1% value at risk (VaR) is the quantity such that the probability of having a loss larger than this VaR tomorrow is 1%. VaR is recommended in the Basel II framework as the basis for capital requirements against market risk. A sound way to estimate VaR would be:

1. Simulate many replications of a vector of risk factors that determine tomorrow's price for the derivative securities.
2. Simulate conditional on each vector of risk factors to estimate the price of all of the derivatives in that scenario, and thus the associated loss.
3. Estimate the VaR by making it the quantity such that in 1% of the scenarios, the portfolio's estimated loss is larger than VaR.

The difficulty is that simulation is required to price many of the derivatives, and one might need to generate, for each of one thousand scenarios, ten thousand paths of one hundred time steps and one hundred state variables, for a total of one hundred billion primitive simulation operations.

Despite advances in computing technology, it is not yet affordable to do this in one night, and consequently financial institutions rely on methodologies of questionable soundness for computing VaR. One of the most well-known is the delta-gamma approximation of the portfolio's value as a quadratic function of the risk factors. The "Greeks" delta and gamma are derivatives of the securities' prices with respect to risk factors. The approximation tends to be accurate only locally, which is unfortunate because large losses (with which VaR is concerned) usually involve large moves in financial markets.

Furthermore, one might hope to use firmwide risk measurement to guide hedging. At present, hedging is conducted primarily by individual traders and usually on the basis of Greeks, in part because of the difficulty of computing an optimal hedge for a large portfolio. There is an opportunity for better and cheaper hedging. However, it will be even more expensive than to estimate a risk measure to solve a simulation-optimization problem to find the optimal hedge that minimizes the risk measure (or some function involving a risk-reward tradeoff).

For these reasons, it would help risk managers to have very efficient simulation procedures for estimating risk measures. Of course, one of the major benefits of Monte Carlo simulation as a tool for numerical computation is that it tends to provide extremely helpful statistical error estimates. We would like to see simulation procedures that efficiently provide confidence intervals for risk measures.

2 PROBLEM STATEMENT

Monte Carlo simulation is often used to estimate the mean of a random variable. It is not necessary for the simulation analyst to possess an explicit expression for the distribution F_V of the random variable V whose mean is to be estimated. All that is required is to be able to sample from its distribution by sampling some basic random vector Z and evaluating a function that maps it to V. Abusing notation slightly, let $V(\cdot)$ also represent this function: $V = V(Z)$. The mean

$$\mathrm{E}[V] = \int V(z)\, dF_Z(z)$$

where F_Z is the distribution of Z; again, F_Z need not be known, but one must be able to sample from it. This enables one to sample from F_V: when Z is sampled from F_Z, $V(Z) \sim F_V$.

Therefore one can also design simulation procedures to estimate functionals T, other than the mean, of the distribution F_V. Where $V_i = V(Z_i)$ and $Z_1, \ldots, Z_k$ are sampled from F_Z, let $\hat{F}_V$ be the empirical distribution of $V_1, \ldots, V_k$, and $T(\hat{F}_V)$ is a point estimate of $T(F_V)$, the quantity of interest. One can also derive (asymptotically valid) confidence intervals for $T(F_V)$.

Some interesting examples of such functionals are:

- The probability $F_V(v)$ that V is less than some value v. This could be useful in studying the rare event of system failure or in evaluating service level agreements.
- For similar reasons, quantiles are of interest. For value at risk, $-\mathrm{VaR}_p$ is the p-quantile of F_V.
- Tail conditional expectation

  $$\mathrm{TCE}_p := \mathrm{E}[-V | V \leq -\mathrm{VaR}_p]$$

  is used as a risk measure by the insurance industry (Manistre and Hancock 2005).
- Expected shortfall

  $$\mathrm{ES}_p \; := \; -\frac{1}{p}\Big(\mathrm{E}\left[V\mathbf{1}_{\{V \leq -\mathrm{VaR}_p\}}\right] + \mathrm{VaR}_p(p - \Pr[V \leq -\mathrm{VaR}_p])\Big)$$

  is a coherent risk measure (Artzner et al. 1999, Acerbi and Tasche 2002).

In the definition of ES, VaR must involve the lower p-quantile if the quantile is not unique. If F_V is continuous and increasing at its p-quantile, TCE and ES are the same (Acerbi and Tasche 2002).

When $T(F_V)$ is related to a quantile of F_V, the binomial distribution of the number of samples exceeding $T(F_V)$ can be used to construct a confidence interval. Confidence intervals for expected shortfall or tail conditional expectation have been derived by Baysal and Staum (2007), based on empirical likelihood, bootstrapping, or on the work of Manistre and Hancock (2005) which makes use of the influence function concept from robust statistics.

Let's call V the *value* of interest and Z the *risk factor*. Examples:

1. Z contains interarrival and service times in a queueing system, and $V(Z)$ is the waiting time of the nth job.
2. Z is a vector of parameters such as interarrival and service rates for a queueing system, whose distribution arises from input modeling, e.g., as a Bayesian posterior distribution. $V(Z)$ is the mean of the stationary distribution of jobs in the system.
3. Z contains financial variables such as stock prices and interest rates at a future date T, and $V(Z)$ is the gain on a financial portfolio over the time interval $[0, T]$.

4. Z describes the characteristics of a disaster, and $V(Z)$ is the expected loss sustained after a response to mitigate the disaster is chosen once Z has been observed.

In the first example, $V(Z)$ can be evaluated: it is a matter of the logic describing the flow of jobs in the queueing system. In the second example, $V(Z)$ can usually not be evaluated directly; it would have to be estimated by simulation of the queueing system, with sampling done conditional on Z. In the third example, if the portfolio contains only securities whose prices at time T are known as a function of Z, then $V(Z)$ can be evaluated; usually, it has to be estimated by simulation of the discounted payoffs of the securities, conditional on the risk factor Z. In the fourth example, $V(Z)$ usually has to be estimated by simulation optimization.

3 WHEN AND HOW TO DO TWO-LEVEL SIMULATION

We now focus on situations where $V(Z) = \mathrm{E}[X|Z]$, and it is possible to sample from the conditional distribution $F_{X|Z=z}$ of X given $Z = z$. To estimate a functional $T(F_V)$, one can use a two-level simulation. At the *outer* level, risk factors $Z_1, \ldots, Z_k$ are sampled from F_Z. At the *inner* level, for each $i = 1, \ldots, k$, *payoffs* $X_{i1}, \ldots, X_{iN_i}$ are sampled from $F_{X|Z=Z_i}$, and one may estimate $V_i = V(Z_i)$ by $\bar{X}_i := \sum_{j=1}^{N_i} X_{ij}/N_i$. In the financial context, the total cost of sampling risk factors is usually negligible compared to the total cost of sampling payoffs, so we regard the computational cost of the two-level simulation as $\sum_{i=1}^{k} N_i$. Let $\hat{F}_{\bar{X}}$ be the empirical distribution of $\bar{X}_1, \ldots, \bar{X}_k$. Then $T(\hat{F}_{\bar{X}})$ is a point estimate of $T(F_V)$.

If the functional of interest T is the mean, a simplification occurs: the mean

$$T(F_V) = \mathrm{E}[V] = \mathrm{E}[\mathrm{E}[X|Z]] = \mathrm{E}[X]$$

is

$$\int v\, dF_V(v) = \int \int x\, dF_{X|Z=z}(x)\, dF_Z(z).$$

In this case, the two levels collapse, and it suffices to do an ordinary one-level simulation in which one samples X by first sampling Z from F_Z, then X from $F_{X|Z}$.

If the function $V(\cdot)$ is known, then only the outer level of simulation is required.

If F_Z is a discrete distribution over a small number k of outcomes whose probabilities are known, then only the inner level of simulation is required. An example of such a situation is estimation of the value of the best of k systems. This is related to ranking and selection, but the difference is that here we are interested in estimating the value of the best system, not finding the identity of the best system or estimating the value of the selected system. This special case fits into our framework when F_Z is a uniform distribution over k systems. Hence in the absence of ties, F_V is a uniform distribution over k values, the functional T finds the distribution's maximum value, and stratification is used in sampling from F_Z. This problem is studied by Lesnevski et al. (2005, 2006), who develop procedures which, applied to risk management, allow for the efficient simulation of coherent risk measures based on generalized scenarios.

If T is not the mean, the function $V(\cdot)$ is unknown, and the statistical error associated with sampling from F_Z can not be entirely removed, then a two-level simulation is required. In what remains, we assume that this is so and we focus on the third example of a financial portfolio. For this example, X is the sum of the discounted payoffs of the securities in the portfolio. To evaluate such a payoff given the risk factor Z representing the market at time T, one must simulate the market's evolution from time T until the security's maturity, when it makes its last payment.

4 CONFIDENCE INTERVALS

How does statistical uncertainty at the inner level combine with statistical uncertainty at the outer level? This question must be answered to create a confidence interval after a two-level simulation.

One sign of the difficulties is that uncertainty at the inner level can lead to bias in the obvious point estimates of $T(F_V)$. This was seen already in the work of Lesnevski et al. (2005), in which there was no uncertainty at the outer level (hence $T(F_V) = T(\hat{F}_V)$). Let $V_{[1]}, \ldots, V_{[k]}$ represent the (increasing) order statistics of the sample $V_1, \ldots, V_k$. For TCE_p, where kp is an integer, $T(\hat{F}_V) = -\sum_{i=1}^{kp} V_{[i]}/kp$. The inner level provides an estimate $\bar{X}_i$ of each V_i, but it does not reveal with certainty the identity of $[i]$, that is, which risk factor results in the ith smallest value. Where $\bar{X}_{(1)}, \ldots, \bar{X}_{(k)}$ represent the order statistics of the sample $\bar{X}_1, \ldots, \bar{X}_k$, the obvious point estimate of $T(\hat{F}_V)$ is $-\sum_{i=1}^{kp} \bar{X}_{(i)}/kp$. Its expectation is $-\sum_{i=1}^{kp} V_{(i)}/kp < -\sum_{i=1}^{kp} V_{[i]}/kp$, so it is a biased estimate of $T(\hat{F}_V)$.

The problem can be viewed as follows. If $V(\cdot)$ were known, a confidence interval $[L(\mathbf{V}), U(\mathbf{V})]$ for $T(F_V)$ could be computed from the sample $\mathbf{V} := (V_1, \ldots, V_k)$. In simulation applications, we are generally unable to place F_V in a parametric family and use parametric hypothesis testing to construct a confidence interval with small-sample validity, but the sample size k is large. Therefore we rely on nonparametric methods to construct a confidence interval satisfying

$$\lim_{k\to\infty} \Pr\{T(F_V) \in [L(\mathbf{V}), U(\mathbf{V})]\} = 1 - \alpha_o. \quad (1)$$

For example, in the ordinary case where T is the mean, the central limit theorem justifies a confidence interval based on asymptotic normality of the sample average of $T(\hat{F}_V)$; it requires computing the sample average and sample standard deviation. For other suitable T, including TCE (Manistre and Hancock 2005), computing what is called the "influence function" of T in the theory of robust statistics justifies a confidence interval centered on $T(\hat{F}_V)$, whose width is proportional to an estimate of the asymptotic standard deviation of $T(\hat{F}_V)$. Alternative nonparametric methods include bootstrapping and empirical likelihood (Baysal and Staum 2007). Because $V(\cdot)$ is not known, $\mathbf{V}$ and hence $L(\mathbf{V})$ and $U(\mathbf{V})$ are unobservable.

Instead we construct $\hat{L}$ and $\hat{U}$ from the observable sample of payoffs to satisfy, for all $\mathbf{V}$,

$$\lim_{\mathbf{N}\to\infty} \Pr\left\{[L(\mathbf{V}), U(\mathbf{V})] \subseteq \left[\hat{L}, \hat{U}\right]\right\} \geq 1 - \alpha_i \quad (2)$$

where $\mathbf{N} := (N_1, \ldots, N_k)$. Applying the Bonferroni inequality to Inequalities (1) and (2),

$$\lim_{k,\mathbf{N}\to\infty} \Pr\left\{T(F_V) \in \left[\hat{L}, \hat{U}\right]\right\} \geq 1 - \alpha \quad (3)$$

where $\alpha = \alpha_o + \alpha_i$, an error spending decomposition which expresses the fact that error arises in one of two ways: the unknown outer-level confidence interval $[L(\mathbf{V}), U(\mathbf{V})]$ fails to cover the true value due to unlucky sampling of risk factors, or the known confidence interval $[\hat{L}, \hat{U}]$ fails to cover the unknown outer-level confidence interval due to unlucky sampling of payoffs at the inner level. The way to construct $\hat{L}$ and $\hat{U}$ so that Inequality (2) can be established depends on the functions L and U providing the outer-level confidence interval, which in turn depend on the functional T and the method of justifying the outer-level confidence interval.

Suppose that sampling is done so the risk factors $Z_1, \ldots, Z_k$ are independent and $X_{i_1 j_1}, X_{i_2 j_2}$ are independent of each other conditional on Z_{i_1} and Z_{i_2} for $i_1 \neq i_2$. As $\mathbf{N} \to \infty$, $\bar{\mathbf{X}}$ converges to multivariate normal with independent components. Each $(\bar{X}_i - V_i)/\sqrt{N_i}$ converges to a normal random variable with mean 0 and variance which can be estimated by the sample variance S_i^2 of $X_{i1}, \ldots, X_{iN_i}$. Define $\epsilon := 1 - (1 - \alpha_i)^{1/k}$ and let t_i be the quantile at the $1 - \epsilon/2$ level of the Student t distribution with $N_i - 1$ degrees of freedom. Let $\mathcal{V}$ be the random k-dimensional box formed as the Cartesian product over $i = 1, \ldots, k$ of the intervals

$$\left[\bar{X}_i - \frac{t_i S_i}{\sqrt{N_i}}, \bar{X}_i + \frac{t_i S_i}{\sqrt{N_i}}\right]. \quad (4)$$

Then $\mathcal{V}$ is an asymptotically valid confidence region for $\mathbf{V}$:

$$\lim_{\mathbf{N}\to\infty} \Pr\{\mathbf{V} \in \mathcal{V}\} \geq (1 - \epsilon)^k = 1 - \alpha_i.$$

Consequently, Inequality (2) is satisfied with

$$\hat{L} = \inf_{v\in\mathcal{V}} L(v) \quad \text{and} \quad \hat{U} = \sup_{v\in\mathcal{V}} U(v). \quad (5)$$

It would also be possible to construct $\mathcal{V}$ in different ways, e.g., as an ellipsoid based on the limiting χ_k^2 distribution of $\frac{1}{k}\sum_{i=1}^k (\bar{X}_i - V_i)^2/S_i^2 N_i$. The ease of performing the optimizations in Equation (5) and the size of the resulting confidence interval are relevant considerations in choosing how to construct $\mathcal{V}$. These optimizations are trivial when L and U are monotone and $\mathcal{V}$ has a least and a greatest element with respect to the usual partial ordering on $\Re^k$, as it does when defined as the k-dimensional box.

For example, to use empirical likelihood (Owen 2001) at the outer level, let

$$\mathcal{S} := \left\{w \in \mathcal{W} \middle| \prod_{i=1}^k (k w_i) \geq \exp\left(-\frac{1}{2}\chi^2_{1,1-\alpha_o}\right)\right\}$$

where $\mathcal{W} := \{w \in \Re^k | w \geq 0, \sum_{i=1}^k w_i = 1\}$. Define $\hat{T}(w, v) := T(F_{w,v})$ where $F_{w,v}$ is the discrete distribution putting probability mass w_i on the value v_i. Then $L(\mathbf{V}) = \inf_{w\in\mathcal{S}} \hat{T}(w, \mathbf{V})$ and $U(\mathbf{V}) = \sup_{w\in\mathcal{S}} \hat{T}(w, \mathbf{V})$. If T is monotone (as is true of mean, quantile, TCE, and ES), then L and U are monotone.

We have constructed a procedure that produces a confidence interval for TCE based on empirical likelihood and proved that it satisfies Inequalities (1) and (2). This procedure uses an inner-level confidence region $\mathcal{V}$ tailored to TCE, different from the box and ellipsoid mentioned above.

5 EFFICIENCY ENHANCEMENT

Efficiency is important for two-level simulation, which can be very computationally expensive, as discussed in Section 1. Efficiency enhancement for two-level simulation is interesting for several reasons. It may not be straightforward to get a confidence interval by the methods described in Section 4 when using variance reduction. The shape of the inner-level confidence region $\mathcal{V}$ greatly influences efficiency. There are questions about how to allocate computational resources between the outer and inner levels, and how to allocate computational resources within the inner level.

5.1 Confidence Region Construction

A good choice of inner-level confidence region $\mathcal{V}$ depends on the risk measure that is being estimated and the method of constructing an outer-level confidence interval, that is, on the functions L and U that provide the lower and upper confidence limits in Equation (1). The construction of the two-level confidence interval as specified in Equation (5) tends to be too conservative and have an excessive coverage probability: Inequality (3) is not tight. The same confidence interval $[\inf_{v\in\mathcal{V}} L(v), \sup_{v\in\mathcal{V}} U(v)]$ can be generated by more than one inner-level confidence region. The largest $\mathcal{V}'$ satisfying

$$\left[\inf_{v\in\mathcal{V}} L(v), \sup_{v\in\mathcal{V}} U(v)\right] = \left[\inf_{v\in\mathcal{V}'} L(v), \sup_{v\in\mathcal{V}'} U(v)\right]$$

is

$$\mathcal{V}' = \left\{v' \mid L(v') \geq \inf_{v\in\mathcal{V}} L(v), U(v') \leq \sup_{v\in\mathcal{V}} U(v)\right\},$$

which can have $\Pr\{\mathbf{V} \in \mathcal{V}'\}$ significantly larger than $1 - \alpha_i$. To increase efficiency, choose $\mathcal{V}$ such that $[\inf_{v\in\mathcal{V}} L(v), \sup_{v\in\mathcal{V}} U(v)]$ is narrow and $\Pr\{\mathbf{V} \in \mathcal{V}'\}$ is near $1 - \alpha_i$.

5.2 Computational Resource Allocation

Given a fixed computational budget C, what number k of risk factors minimizes the width $\hat{U} - \hat{L}$ of the confidence interval? Roughly, we suppose that $\sum_{i=1}^{k} N_k = C$: nearly all of the cost is inner-level simulation. Lee (1998) answers a related question: he takes $N_1 = \cdots = N_k = N$ and finds that $k \propto C^{2/3}$ minimizes the asymptotic MSE of an estimate of VaR. The answer to our question may differ, and it depends on the risk measure and on how the confidence interval $[\hat{L}, \hat{U}]$ is constructed.

Given k, how should $N_1, \ldots, N_k$ be chosen? Lesnevski et al. (2005, 2006) provide several techniques for increasing the efficiency of the inner-level simulation when $T(\hat{F}_V) = V_{[k]}$.

One technique that is widely applicable is the use of multi-stage simulation at the inner level. After a first stage in which n_0 observations are simulated conditional on each of the risk factors, one can set N_i proportional to a sample variance $S_i^2(n_0)$. This can help when T is TCE: much as in stratified sampling, when N_i is not proportional to $\mathrm{Var}[X|Z_i]$, there is a difference in the marginal benefit of sampling payoffs conditional on different risk factors.

Screening is another technique for increasing efficiency. It is applicable to TCE or whenever $L(\mathbf{V})$ and $U(\mathbf{V})$ do not depend on all of the order statistics $V_{[1]}, \ldots, V_{[k]}$. For TCE, after the first stage, those risk factors whose sample means $\sum_{j=1}^{n_0} X_{ij}/n_0$ are too large compared to those of other risk factors are "screened out." Let $I \subseteq \{1, \ldots, k\}$ contain the indices of those that are not screened out. Only for $i \in I$ do we then simulate $X_{i,n_0+1}, \ldots, X_{i,N_i}$. This creates a great savings, but at a cost, familiar from ranking and selection procedures that use screening. The error spending structure must be modified: $\alpha_i = \alpha'_I + \alpha_S$, where α_S is spent on the possibility of an error in screening, while α'_i must replace α_i in Inequality (2). The details are not as straightforward as in Lesnevski et al. (2005) where the goal was to get $[k] \in I$. However, we have constructed a procedure that satisfies this modified error spending structure for TCE and empirical likelihood, where correct screening means $\{[1], \ldots, [\ell]\} \subseteq I$.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. DMI-0555485.

REFERENCES

Acerbi, C., and D. Tasche. 2002. On the coherence of expected shortfall. *Journal of Banking and Finance* 26: 1487–1503.

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath. 1999. Coherent measures of risk. *Mathematical Finance* 9: 203–228.

Bank for International Settlements. 2004. Basel II: International convergence of capital measurement and capital standards: a revised framework.

Baysal, E., and J. Staum. 2007. Empirical likelihood for value at risk and expected shortfall. Working Paper 07-01, Dept. of Industrial Engineering and Management Sciences, Northwestern University.

Lee, S.-H. 1998. Monte Carlo computation of conditional expectation quantiles. Dissertation, Dept. of Operations Research, Stanford University.

Lesnevski, V., B. L. Nelson, and J. Staum. 2005. Simulation of coherent risk measures based on generalized scenarios. Forthcoming, *Management Science*.

Lesnevski, V., B. L. Nelson, and J. Staum. 2006. An adaptive procedure for estimating coherent risk measures based on generalized scenarios. Working Paper 06-05, Dept. of Industrial Engineering and Management Sciences, Northwestern University.

Manistre, B. J., and G. H. Hancock. 2005. Variance of the CTE estimator. *North American Actuarial Journal* 9: 129–154.

Owen, A. B. 2001. *Empirical Likelihood*. New York: Chapman & Hall/CRC.

AUTHOR BIOGRAPHIES

HAI LAN is a Ph. D. student in the Department of Industrial Engineering and Management Sciences at Northwestern University. His email address is <h-lan@northwestern.edu>.

BARRY L. NELSON is the Charles Deering McCormick Professor of Industrial Engineering and Management Sciences at Northwestern University, and is Director of the Master of Engineering Management Program there. His research centers on the design and analysis of computer simulation experiments on models of stochastic systems, and he is Editor in Chief of *Naval Research Logistics*. Nelson has held many positions for the Winter Simulation Conference, including Program Chair in 1997 and current membership on the Board of Directors. His e-mail and web addresses are <nelsonb@northwestern.edu> and <www.iems.northwestern.edu/~nelsonb/>.

JEREMY STAUM is Assistant Professor of Industrial Engineering and Management Sciences at Northwestern University. His research interests center on risk management and simulation in financial engineering. Staum is Associate Editor of *ACM Transactions on Modeling and Computer Simulation*, *Naval Research Logistics*, and *Operations Research*, and is coordinating the Risk Analysis track at the 2007 Winter Simulation Conference. His e-mail address is <j-staum@northwestern.edu> and his web page is <www.iems.northwestern.edu/~staum>.