

ORMS

TODAY

February 2006
Volume 33 • Number 1

PROBABILITY MANAGEMENT

Who will wear the
Chief Probability Officer
hat in your organization?

ALSO INSIDE:

Academe & Industry:

- Can consortia break down the barriers?
- Problem Exchange aims to 'keep it real'
- IPC readies students for real world
- Management class teams up for success


This article appeared in the February 2006 issue of OR/MS Today. It is reproduced here with permission. It may not be reprinted or re-posted on the Web without permission of OR/MS Today.

Planning for an uncertain future calls for a shift in information management – from single numbers to probability distributions – in order to correct the “flaw of averages.” This, in turn, gives rise to the prospect of a Chief Probability Officer to manage the distributions that underlie risk, real portfolios, real options and many other activities in the global economy.

BY SAM SAVAGE,
STEFAN SCHOLTES AND
DANIEL ZWEIDLER



PROBABILITY MANAGEMENT



Today's world economy is driven by global uncertainties such as exchange rates, political upheaval and energy prices, layered upon local uncertainties involving individual projects. Pharmaceutical firms must manage their R&D operations in light of changing regulations and global pandemics on the one hand and uncertain outcomes surrounding specific compounds on the other. Banks must choose their loan portfolios in the face of unpredictable interest rates and global economic factors as well as uncertain regional demographics and competition. Petroleum firms must allocate their exploration budgets across diverse geographical regions and new technologies, given global uncertainties in oil price and geopolitics and local uncertainties concerning geology and markets.

These uncertainties create an unprecedented number of interdependent risks. Modern financial theory recognizes that economic return entails such risk [1,2]. Further, it tells us that the risk of a portfolio of investments is not merely an additive property of the individual investments, but is driven by their interdependence. If the underlying statistical relationships of these uncertainties are captured in the planning process, they can be exploited to find optimal risk-based tradeoffs between strategic objectives. If they are ignored, large risks may be masked and significant mitigation and economic return opportunities will remain untapped.

This perspective is nearly universal among managers of portfolios of securities, and statistical relationships are arguably even more important in *real portfolios*. (We use the term "real portfolio" for a portfolio of projects rather than financial instruments in the same way the term "real option" is used for options involving projects rather than financial assets.) Unfortunately, most organizations lack a consistent approach to modeling and communicating the underlying statistical relationships between business units. Instead, they typically use single average or base-case numbers to represent uncertain business parameters and metrics. This leads to a class of systematic errors known as the *flaw of averages* [3].

The authors encourage an area of management focus, often ignored today, that can correct the flaw of averages. What is needed is a shift in information management, from single numbers to probability distributions. We call this area *probability management*, and argue that it is a prerequisite for the effective management of risk, real portfolios, real options and many other activities in the global economy. In this article we begin our discussion in broad terms, using an analogy with the incandescent light bulb and electric power grid. We then revisit some of the tacit assumptions of business planning under uncertainty, highlighting the flaw of averages and the seven deadly sins of averaging. Next, we describe an approach to probability management, developed by the authors, that we call *coherent modeling*. We then outline our ongoing experience in applying these ideas to

the planning cycle within a major petroleum company. Finally, we make a brief comparison of probability management with the current practice of risk management.

An Analogy

BY 1880, THOMAS EDISON had developed a good incandescent light bulb [4]. However, the market for this invention was small, as it was of no practical value without a source of electricity. To actually get light from a bulb required the purchase of an expensive generator and knowledge of electrical theory. The first modern transmission of alternating current based on the theories of Nikola Tesla [5] did not occur for another decade. With standardized sources of electricity, neither generators nor theoretical knowledge were required of the end user, and the market for light bulbs and other appliances exploded.

Today, simulation does for uncertainty what the light bulb of 1880 did for darkness. (We use the word "simulation" loosely to mean any sort of stochastic analysis based on modeling probability distributions through sampling.) If properly used, it can illuminate. Simulations, however, require probability distributions for their uncertain inputs, much as light bulbs require electricity. Currently, users of simulation need to specify the type of distributions used to generate their input values. This is analogous to requiring the users of light bulbs to generate their own electricity.

Probability management is based on three underpinnings, which we will describe in terms of this analogy: 1. interactive simulation, 2. stochastic libraries, and 3. certification authority.

Interactive simulation tools play the role of light bulbs by illuminating uncertainty and risk for a wide population of managers. New technologies will run simulations nearly instantaneously each time the parameter of a business model is changed. Interactive visual feedback will provide management with an experiential understanding of uncertainty and risk [6].

Stochastic libraries contain certified probability distributions for use in simulations throughout an organization. They are analogous to the electric power grid. By providing a ready source of input distributions in standardized formats, both theoretical knowledge and effort on the part of the end user are greatly reduced, facilitating the use of probabilistic modeling.

Certification authority is required for the distributions in the stochastic libraries of an organization much in the way the local power authority ensures that you always get a standard voltage from your wall socket. A suggested name for this certifying authority is the *Chief Probability Officer* (CPO), and the person or office wearing this hat requires a combination of both statistical and managerial skills. Ultimately the CPO must find the right balance between authorizing complex multivariate statistical time series,

which only a few specialists understand, versus single “average” scenarios, leading to the flaw of averages.

Business Planning Under Uncertainty

RECENT GOSSIP IN THE exploration and production departments of petroleum companies suggests that E&P stands for “Excel and PowerPoint.” While the endemic use of PowerPoint slides for communicating technical data is problematic, as recognized by the Columbia Accident Investigation Board [7], it is hard to imagine planning without spreadsheet models. We suggest Michael Schrage’s book “Serious Play” [8] for an in-depth account of how spreadsheets allow managers to quickly prototype alternate models of their enterprises.

There is, however, one area in which spreadsheet models badly miss the mark. They inadequately account for uncertainty and risk. Future projections of metrics such as demand, prices and costs are often condensed into a single “average” or “base case” value, which serves as input to the model. The resulting performance metrics are then expressed as single “average” outputs. The justification is that “the scenario we use in the model is our best estimate,” implying tacitly that the resulting outputs are also the best estimate of performance. This results in a variety of systematic errors, which, although documented in probability textbooks for decades, are rarely recognized in practice. Collectively we call these errors the “flaw of averages.”

Before describing the flaw of averages it is useful to distinguish between uncertainty and risk. Although the literature presents numerous definitions, the authors prefer the following ones, which are consistent with the theories of probability and utility:

Uncertainty is an objective feature of the universe over which you have no control. Uncertain quantities such as the weather, the card you draw from a shuffled deck and tomorrow’s price of gold are what mathematicians call *random variables*. The best you can do to estimate a random variable *a priori*, is to estimate its probability distribution. These are the uncertain inputs to a model.

Risk is in the eye of the beholder. If I own gold, the risk for me is that gold prices will drop. If I have shorted gold, the risk for me is that prices will rise. From the authors’ perspective, risk involves a formula fed by one or more random variables. This is known by mathematicians as a *function of random variables*, or by spreadsheet users, as a formula with uncertain inputs. These correspond to the output metrics of our business model.

We will review these concepts in terms of a sobering example of the flaw of averages. Consider a drunk staggering down the middle of a busy highway. The position of the drunk is a random input to the model, with an average of the centerline. The output metric of interest is the physical state of the drunk. A prediction of the future state of the

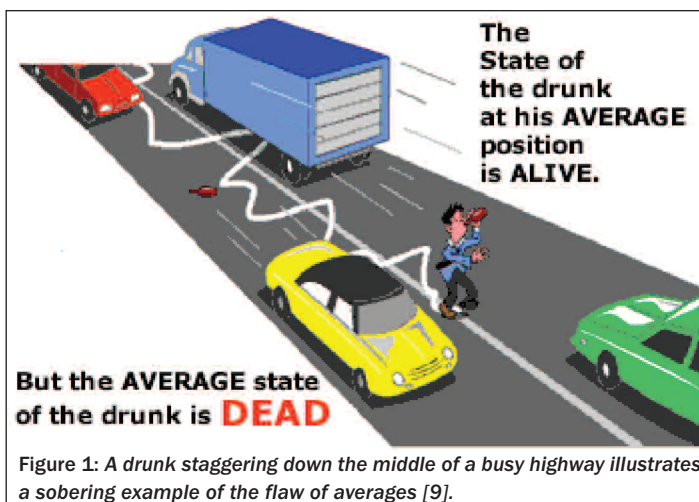


Figure 1: A drunk staggering down the middle of a busy highway illustrates a sobering example of the flaw of averages [9].

drunk based on his *average* position will claim that he is alive. However, on *average* he is clearly dead (Figure 1).

With the above example in mind, we present several other forms of the flaw of averages:


The Seven Deadly Sins of Averaging

1. The Family with 1 1/2 Children
2. Why Everything is Behind Schedule
3. The Egg Basket
4. The Risk of Ranking
5. Ignoring Restrictions
6. Ignoring Optionality
7. The Double Whammy

1. **The Family with 1 1/2 Children:** Often the “average” scenario, like the “average” family with 1 1/2 children, is non-existent. For example, a bank may have two main groups of young customers – students with an average income of \$10,000 and young professionals with an average income of \$70,000. Would it make sense for the bank to design products or services for customers with the average income of \$40,000?

2. **Why Everything is Behind Schedule:** Imagine a software project that requires 10 separate subroutines to be developed in parallel. The time to complete each subroutine is uncertain and independent, but known to average three months, with a 50 percent chance of being over or under. It is tempting to estimate the average completion time of the entire project as three months. But for the project to come at three months or less, each of the 10 subroutines must be completed at or below its average duration. The chance of this is the same as flipping 10 sequential heads with a fair coin, or less than one in a thousand!

3. **The Egg Basket:** Consider putting 10 eggs all in the same basket, versus one by one in separate baskets. If there is a 10-percent chance of dropping any particular basket, then either strategy results in an average of nine unbroken eggs. However, the first strategy has a 10-percent chance of losing all the eggs, while with



the second, there is only one chance in 10 billion of losing all the eggs.

4. **The Risk of Ranking:** It is common when choosing a portfolio of capital investment projects to rank them from best to worst, then start at the top of the list and go down until the budget has been exhausted. This flies in the face of modern portfolio theory, which is based on the interdependence of investments. According to the ranking rule, fire insurance is a ridiculous investment because on average it loses money. But insurance doesn't look so bad if you have a house in your portfolio to go along with it.
5. **Ignoring Restrictions:** Consider a capital investment in infrastructure sufficient to provide capacity equal to the "average" of uncertain future demand. It is common to assume that the profit associated with average demand is the average profit. This is generally false. If actual demand is less than average, clearly profit will drop. But if demand is greater than average, the sales are *restricted* by capacity. Thus, there is a downside without an associated upside, and the average profit is less than the profit associated with the average demand.
6. **Ignoring Optionality:** Consider a petroleum property with known marginal production costs and an uncertain future oil price. It is common to value such a property based on the "average" oil price. If oil price is above average, the property is worth a good deal more. But if the price drops below the marginal cost of production, the owners have the *option* to halt production. Thus, there is an upside without an associated downside, and the average value is greater than the value associated with the average oil price. Note that the SEC currently values petroleum properties based on the oil price on Dec. 31 of the preceding year, a clear commission of the flaw of averages [10].
7. **The Double Whammy:** Consider a perishable inventory of goods with uncertain demand, in which the quantity stocked is the "average" demand. If demand exactly equals its average, then there are no costs associated with managing the inventory. However, if demand is less than average then there will be spoilage costs, and if demand is greater than average there will be lost sales costs. So the cost associate with average demand is zero, but average cost is positive.

The seven deadly sins of averaging are by no means exhaustive. However, they are widespread, easy to understand when explained and have serious consequences if ignored.

Coherent Modeling

PROBABILITY MANAGEMENT, LIKE universal peace and happiness, sounds good in principle, but the devil is in the details. The authors have developed a relatively simple approach to probability management, called coherent modeling, which has been found useful in practice. The fundamental idea is a stochastic library consisting of pre-generated

random trials, a throwback to the random number tables of the 1950s [11].

The benefits of the coherent modeling approach are:

1. Statistical dependence is modeled consistently across entire organizations.
2. Probabilistic models may be rolled up between levels of an organization.
3. Probabilistic results may be audited at a later date.

A simple example of coherent modeling is presented in “Stochastic Library Structure” (see box) and will be discussed in more detail in a subsequent article in *OR/MS Today*.

Case Study

IN THE FOLLOWING CASE STUDY, the authors applied coherent modeling to provide a global perspective for a petroleum exploration firm that had traditionally been highly decentralized.

BACKGROUND

Shell Exploration and Production is engaged in the upstream activities of acquiring, exploiting, developing and producing oil and gas. In 2003, Shell reorganized its petroleum exploration according to a global operating model. This meant moving from a highly decentralized business with regional allegiances and reward systems to a single centralized organization managing a large portfolio of exploration opportunities.

Most of the senior staff, despite a large expatriate base, had never worked in a truly global environment. Staff was transferred from one operation unit to another, transferring their knowledge acquired during the previous assignment, but immediately pledging allegiance to the new regional management structure. Rather than a central command supervising a global port-



Figure 2: Management visualizes relationships for various investment levels through group interaction with the model.

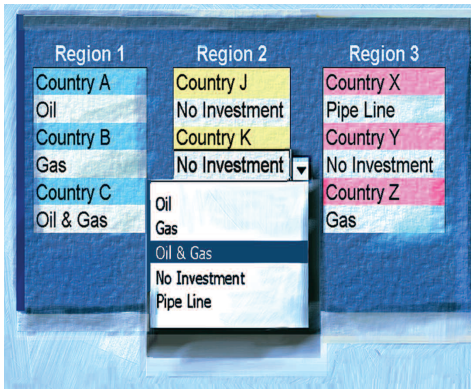


Figure 3: The simple menu interface was a hit with senior management.

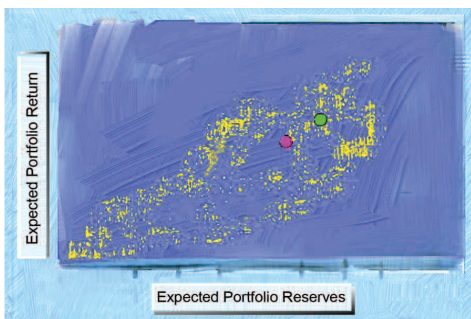


Figure 4: Comparing two particularly portfolios (green and pink dots) in a universe of portfolios.

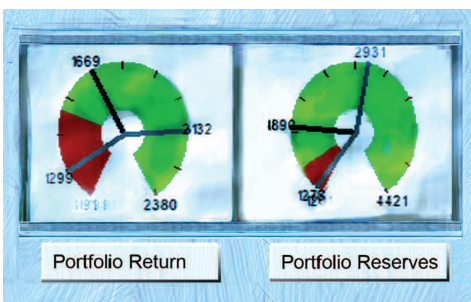


Figure 5: Gauging the probability of falling short of aspirations.

folio of opportunities, there were a large number of entities competing for a limited pot of exploration funds during the annual capital allocation. The methodology used ranked the various opportunities and funded the highest-ranking ones until the money ran out. This approach was further limited by the fact that local imperatives had to be honored with some low-ranking opportunities funded due to real or perceived local commitments.

GLOBAL PERSPECTIVE

After the initial effort of assembling a global portfolio, it quickly became apparent that a new approach to capital allocation was necessary. Rather than a bottom-up assembly of an exploration business plan based on individual opportunities, a more strategic top-down approach had to be designed. In fact, a consequence of the flaw of averages is that the metrics associated with a portfolio of exploration ventures are not merely the sum of the corresponding metrics of the ventures contained in the portfolio. The fundamental business question is: What portfolio of funded ventures is optimally aligned with the overall exploration strategy? The basic approach was to extend the concepts of modern financial portfolio theory developed in the 1950s and 1960s to portfolios of risky exploration venture [12].

STOCHASTIC LIBRARY

A major hurdle was to create a stochastic library based on key venture metrics simple enough to implement, but detailed enough to be credible. Distributions of potential hydrocarbon volumes and economic value, as well as the associated risks, were collected for the various ventures Shell is, or considers, prosecuting. What was fur-

Illustrations by Alice Freund of Corporate Portraits and courtesy of Shell

ther required was an integration of these individual distributions of local uncertainties with global uncertainties such as price and geopolitical events into a library of trials that preserved the statistical relationships between the ventures, that is, was coherent. Individual libraries are created for discrete global scenarios so that the individual impact of a particular scenario on a strategy can be accessed.

Despite the many similarities, when attempting to optimize a portfolio of exploration ventures rather than stocks, a few key differences emerge. Unlike stock portfolios, in which any mix of assets is possible, typically an individual exploration venture must either be in or out of the portfolio. Unlike assets with a market history, there is no direct way to measure the statistical dependence between potential projects; instead one must rely on structural econometric models that relate the projects to each other. When optimizing portfolios of ventures, there is not a simple unique risk/reward tradeoff curve; instead there are many potential tradeoffs between pairs of metrics, reserves vs. revenue, short-term vs. long-term benefits, etc. Thus, a primary goal of the model was to help management visualize these relationships for various investment levels, through group interaction with the model (Figure 2).


The “exploration cockpit” developed for this purpose comprises a limited set of controls and displays. The ventures and their various execution alternatives constituting an individual portfolio are selected or displayed in a standard pull-down menu (Figure 3). This menu interface was so simple and interactive, that senior management was eager to use it in the midst of heated discussion.

The selected portfolio is then highlighted in the cloud of all feasible venture portfolios (Figure 4). Here the green and pink dots repre-

STOCHASTIC LIBRARY STRUCTURE

The simplest element of a coherent stochastic library is a stochastic information packet, or SIP, comprised of a list of trials of some uncertain parameter or metric. For example, consider a petroleum engineer modeling the economic output of an exploration venture at site A. He could generate a SIP of this metric by running and saving 10,000 Monte Carlo trials of a random quantity of oil multiplied by a random price of oil (Figure A).

Figure A




Similarly, other engineers could generate SIPs of their own ventures. At first you might think that the SIPs of the various ventures could all be stored together to form the stochastic library, but this would not be coherent. First, unless each engineer used the same distribution for the price of oil, the results from the various ventures would not be comparable.

Second, even if they used the same distribution for the price of oil, a given trial of the SIP of one venture might have a randomly generated high price of oil while the same trial of the SIP of another venture might have a low price of oil. In reality, the price of oil, although uncertain, is nearly the same world wide, and creates a strong statistical relationship between ventures, which must be preserved.

To make the library coherent, the CPO would make available a certified SIP of the oil price distribution (Figure B).

Figure B



Now, each engineer can run their own Monte Carlo simulation of random oil quantity times oil price, but this time the price values would be drawn sequentially from the certified SIP in their original order. In this case, SIPs of the various ventures would be coherent, and form what we refer to as a stochastic library unit with relationships preserved, or a SLURP (Figure C).

In mathematical terms, the SLURP is a set of samples from a multivariate distribution.

Now suppose the firm wished to know the distribution of total economic return for the entire portfolio of ventures. Merely summing the elements of the SLURP trial by trial across ventures would result in the SIP of the entire portfolio, reflecting the statistical dependence. This is, in effect,

the roll up of a stochastic model from the venture level to the corporate level (see Figure D).

Figure C

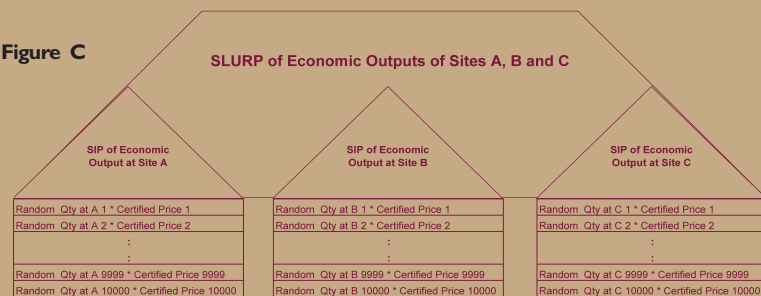
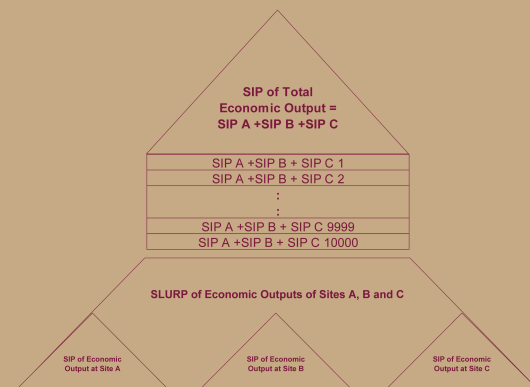


Figure D



sent two portfolios that are being compared, while the yellow cloud displays the limits of the universe of portfolios. In this case, the pair of metrics compared are “expected portfolio reserves” and “expected portfolio return,” but the metrics may be interactively swapped in and out to get different perspectives of the portfolio universe.

Stochastic gauges display adjustable confidence intervals around median values for critical portfolio parameters, highlighting the probabilities of falling short of aspirations defined by the exploration strategy (Figure 5).

These simple but effective displays constituted the sole numerical information available to executives during the top down build of the exploration business plan.

Managerial response: Two sets of workshops were held, the first with the regional planning managers only and the second with the senior exploration executives. It became apparent at the first meeting that the planners were surprised that the points on the graph represented portfolios, rather than individual projects. They quickly grasped this concept and for the first time began to focus on combinations of individual ventures into portfolios and not on individual projects in classical one-dimensional ranking displays. For the first time the question shifted from “How does my venture rank?” to “How does my venture contribute to the portfolio?” Managers who were accustomed to silo thinking were confronted with “Big Picture” issues on the spot.

The acid test was certainly the next workshop with senior executives, who were also not accustomed to looking at portfolios of ventures, albeit they had some prior exposure to the methodology. The same phenomenon was observed, as they were not presented with a direct ranking of exploration projects; they had to shift to a more global perspective. Members of the group now had a source of motivation to operate as a cohesive team in optimizing the overall portfolio. Although there were still obvious temptations for a member to promote their own ventures, thereby increasing their own budgets, the adverse consequences, if any, were now immediately apparent to the entire group.

Does this experience represent the dusk of the decentralized exploration business model? Many challenges remain, including the sustainability of behaviors, the structuring of incentives, quality control of the data, etc. What we can say with certainty is that management gained a perspective into the performance of the venture portfolio as a whole, and that the same approach has been continued for a second year of planning.

Summary

PROBABILITY MANAGEMENT SHIFTS the focus away from trying to predict uncertain future business metrics directly (often the approach in risk management), to understanding the underlying uncertainties that drive those met-

rics. It may be applied within a single business unit, or scaled to model an enterprise, industry or entire economic sector. Some organizations, notably in finance, have been doing it for years. We hope the ideas behind coherent modeling can increase the use of probability management, and help control the flaw of averages in a wide range of organizations.

In a subsequent article, we will describe emerging technologies that are enabling efficient probability management, and point the way to additional areas of application. ◀◀

Sam Savage is a consulting professor of management science and engineering at Stanford University, a collaborator on the spreadsheet optimization package *What's Best*, and founder and president of *AnalyCorp Inc.*, a firm that develops executive education programs and software for improving business analysis.

Stefan Scholtes is a professor of management science and director of research of the Judge Business School at the University of Cambridge. His theoretical research interest in mathematical programming is complemented by applied work that seeks to help managers and engineering designers in their understanding of system values in a complex, uncertain and dynamic environment.

Daniel Zweidler is head of global exploration planning & portfolio for Shell where he helps define the exploration investment case for Shell, merging regional exploration realities and imperatives with new country access opportunities and the competitive landscape. He is responsible for delivering the global exploration and EP growth business plan.

REFERENCES

1. Markowitz, H. M., “Portfolio Selection: Efficient Diversification of Investments,” second edition, Blackwell Publishers, Inc., Malden, Mass. (1957, 1997).
2. Sharpe, William F., “Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk,” *Journal of Finance*, Vol. XIX, No. 3 (September 1964), pgs. 425-442.
3. Savage, Sam L., “The Flaw of Averages,” *Harvard Business Review*, November 2002.
4. <http://www.ideafinder.com/history/inventions/story074.htm>
5. http://en.wikipedia.org/wiki/Electric_power_transmission#History
6. Savage, Sam L., “Blitzograms – Interactive Histograms,” *Informations Transactions on Education*, January 2001 (<http://ite.pubs.informs.org/Vol1No2/Savage/Savage.php>).
7. Report of Columbia Accident Investigation Board, Vol. I, Aug. 26, 2003 (http://www.nasa.gov/columbia/home/CAIB_Vol1.html).
8. Schrage, Michael, “Serious Play: How the World’s Best Companies Simulate to Innovate,” Harvard University Press, 2000.
9. Illustration from Savage, Sam L., “Decision Making with Insight,” Duxbury Press, Belmont Calif., 2003, reproduced with permission of the publisher.
10. Carlisle, T., “How Lowly Bitumen Is Biting Oil Reserve Tallies,” *Wall St. Journal*, Feb.14, 2005.
11. <http://www.rand.org/publications/classics/randomdigits/>
12. Ball, Ben C. and Savage, Sam L., “Holistic vs. Hole-istic Exploration and Production Strategies,” *Journal of Petroleum Technology*, September 1999.