SPEAKING UNCERTAINTY TO POWER
Risk-Aware Forecasting and Budgeting

BY SHAYNE KAVANAGH AND SAM SAVAGE

This article is adapted from the upcoming GFOA publication, Informed Decision-Making through Forecasting: A Practitioners Guide to Government Revenue Analysis, *by Shayne Kavanagh.*

Budget forecasts are almost always wrong; it is just a question of by how much, and in which direction. The ways in which finance officers handle the risk of a forecast being wrong can have a real impact on their credibility and the quality of the budgetary decisions made by the chief executive officer and/or governing board.

Two variables affect the accuracy of a forecast: the inherent uncertainty of tax (or other) revenues, and the extent to which that uncertainty is taken into account by decision makers. Many people think good forecasting is simply a matter of accuracy — that is, making precise predictions that are used to inform the development of a given budget. But in practice, good forecasting also involves an element of education. The likelihood of different scenarios should be conveyed, and the consequences of budgeting mistakes should be transparent.

Two variables affect the accuracy of a forecast: the inherent uncertainty of tax (or other) revenues, and the extent to which that uncertainty is taken into account by decision makers. Many people think good forecasting is simply a matter of accuracy — that is, making precise predictions that are used to inform the development of a given budget.

TRADITIONAL FORECASTING: FLAWS AND CHALLENGES

Forecasts are the starting point for budgeting. Forecasts are typically presented as a single-number "point" estimate. This not only obscures the range of possible revenue outcomes that really exist, but also leads to systematic errors that do not average out in the long run. This is commonly known as the Flaw of Averages (also known as Jensen's Inequality). The single-point forecast also fixates the audience on a single potential outcome, so they may not plan adequately for alternative futures when developing a budget. There are negative consequences if actual revenues miss the point estimate on either side. Therefore, even if your forecasts are right on average, you will be systematically exposed to negative budgetary consequences. The consequences of underestimating revenues are typically less severe than those of overestimating revenues — if revenues are overestimated, the government may have to cancel spending plans, and, in the worst case, lay off staff.

Conservative Forecasts. Because revenue overestimates have the potential to create a budget shortfall and the need for mid-year spending cuts, many finance officers pursue a strategy of conservative revenue estimating — but this conservative forecasting strategy is not without consequence. If revenues are underestimated, the government will have foregone the opportunity to budget for a potentially valuable public service. For instance, a school district might have been able to fund reading help for a greater number of struggling students, or a city might have been able to put more police officers on the streets. Revenue underestimates might also cause unnecessary stress during budget development as decision makers struggle to make trade-offs that are unnecessary.

A conservative forecasting strategy might also have consequences for the finance officer's credibility, hurting his or her ability to provide financial leadership. GFOA's research shows that the strategy can work if policymakers are in agreement with it; otherwise, it can backfire. For example, a city that GFOA studied had a city council that expected revenue forecasts to represent the most objective approximation of what actual revenues would be. Consistent underestimates led the council to accuse the finance officer of "playing games" with the budget.

Best Estimates. Imagine that instead of a conservative forecast, the finance officer provides a "best estimate" of revenues — the single number that the finance officer believes has the best chance of matching actual revenues. This will lead to a smaller error on average, but a greater risk of an actual budget shortfall during the year. Certainly, an actual shortfall won't help the finance officer's credibility either. Furthermore, the finance officer's best estimate will never be 100 percent accurate. When the forecast is presented as single-point estimate, the audience may not plan adequately for futures that differ from that estimate.

Conservative and best-estimate forecasting strategies put the finance officer in a lose-lose situation. Both strategies can also lead to less than optimal resource allocation decisions. A conservative approach is more likely to lead to lost opportunities, whereas the best estimate creates greater exposure to financial stress. When either approach is presented as a single-point forecast, it will not lead the audience to plan adequately for an uncertain future. What can be done?

> Forecasts are typically presented as a single-number "point" estimate. This not only obscures the range of possible revenue outcomes that really exist but also leads to systematic errors that do not average out in the long run.

This article will show how several cities are using spreadsheets to bring sophisticated risk awareness to municipal finance. They demonstrate that finance officers don't need to choose between giving budget decision makers a conservative forecast and a best estimate. Instead, by explicitly recognizing the uncertainty in forecasts, finance officers can improve their dialog with decision makers and arrive at a budget that makes the best use of all available resources, while mitigating the risk posed by revenue shortfalls.

RISK-AWARE FORECASTING

As a result of the Great Recession, the City of Colorado Springs, Colorado, had spent its reserves down to an amount the mayor found unacceptable. The city needed a way to make better decisions about budgeting for current services while replenishing its reserves. The volatility of sales taxes — the city's primary revenue source — made it difficult to plan for these competing goals. The city needed to have a risk-aware discussion of its options so the mayor and the city council could establish the proper balance between current spending and reserve replenishment. Producing a conservative forecast in hopes of producing a large end-of-year surplus could endanger the credibility of the finance staff because they would be denying elected officials the information needed for a fully informed discussion of the options. However, simply providing a point forecast of their best estimate of future revenues would not address the risk posed by the volatility of the sales tax — any plans the elected officials might make based on such a forecast would have been insufficiently informed of an uncertain future.

For the fiscal year 2011 budget, the city finance staff presented a "risk-aware forecast" — a forecast that explicitly acknowledges the risk that a given level of revenue will not be met.

City finance staff started out by making the case that revenue forecasts are best understood as a range of possible future revenues rather than single-point forecasts. This approach gave elected officials a choice. A higher value in the range would allow for more planned expenditures on services, but with a greater risk that the forecast wouldn't be met, therefore potentially depleting the fund balance. A lower value in the range would restrain additional planned expenditures on services, but with a higher chance of the forecast being met or exceeded, therefore potentially adding to fund balance.

Second, finance staff developed a risk-aware presentation. To do so, they assumed that the probability of different revenue outcomes takes the shape of a normal distribution, or bell curve.[1] Given a few years of history, it is easy to calculate a normal distribution in Excel. The process can even be automated.

Exhibit 1 illustrates the Colorado Springs forecast for its revenues as a normal distribution. The middle of the distribution is a single number that the staff believes will be closest to the actual revenues on average. This is the city staff's best estimate of what future revenues will be.

Using the mathematical properties of the normal distribution, the city calculated the odds of actual revenue being

Exhibit 1: Colorado Springs Presentation of a Risk-Aware Forecast

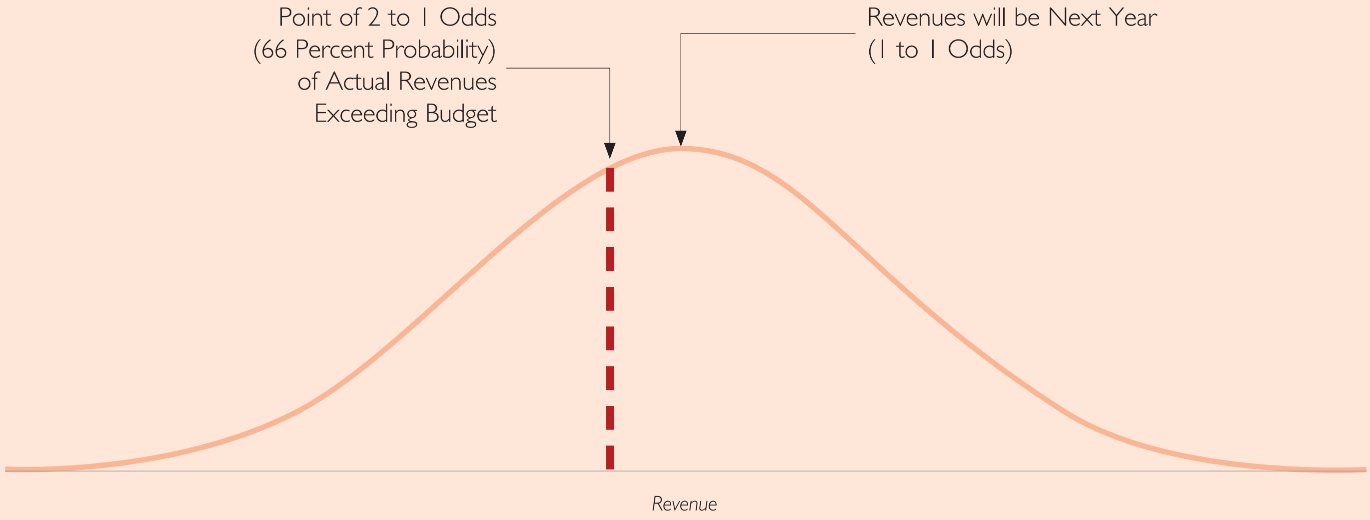

Why Did Colorado Springs Stay on the Left-Hand Side of the Curve?

A common concern finance officers have with a best-estimate forecast is that it would encourage decision makers to spend up to that amount. However, a risk-aware forecast like the one produced by Colorado Springs demonstrates that the best-estimate forecast provides 1:1 odds, or 50 percent probability, of experiencing a revenue shortfall during the year. Psychological research shows that people tend to be loss averse — that is, the prospect of a loss weighs more heavily than the prospect of an equal-sized gain (about twice as heavily, in fact). Hence, most people will not be comfortable staking the financial health of the organization on the same probability as that of a coin flip. When presented with a risk-aware best-estimate forecast, they will likely pick a conservative budgeting strategy.

Once you know the mean and standard deviation of your uncertain revenue, it is easy to calculate the chance and odds of achieving any particular target, as shown in the diagram below.

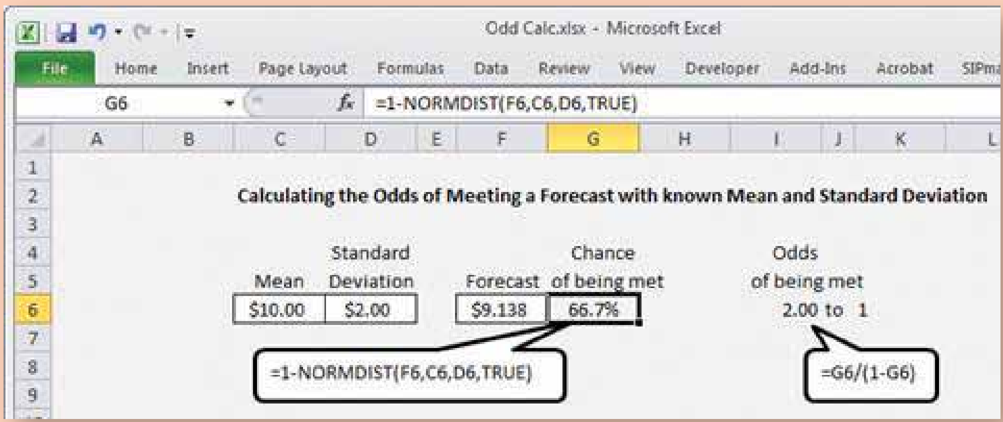

greater than the best-estimate forecast. For example, the odds are 1:1 at the peak of the normal distribution because there is an equal chance of actual revenue being greater than or less than the best-estimate forecast. At the end of the Great Recession, knowing that end-of-year surpluses could be used to replenish the city's partially depleted reserves, the elected officials were invited to select the odds that they would be comfortable with for building a budget. After some discussion, they settled on a budgeted revenue figure that provided 2:1 odds, or a 66 percent chance, that there would be a surplus. Put another way, the elected officials deliberately lowered the amount they would budget in exchange for a greater chance of realizing an end-of-year surplus.

> The consequences of underestimating revenues are typically less severe than those of overestimating revenues — if revenues are overestimated, the government may have to cancel spending plans, and, in the worst case, lay off staff.

The city has gradually built its reserves back since 2011. Over time, elected officials have chosen to reduce the odds somewhat, thereby increasing the amount of money available for current services while still providing a good chance that actual revenue will exceed budget. Each year there is a discussion of the level of risk they are willing to take on, and odds are selected. For example, elected officials selected 1.7 to 1 odds for the 2016 budget.

Besides helping the city replenish its reserves, the risk-aware forecast presentation also helped increase the credibility of the finance office's projections. Before moving to this presentation, elected officials would often "balance" the budget by increasing the forecast — a sure sign that officials do not see the forecast as credible estimate of what future revenues would actually be. After moving to this presentation, officials don't seek to change the estimate; instead, they discuss if they are willing to take on more financial risk in order to fund more current services.

FORECASTING AND BUDGETING

The cities of Redmond, Washington; Boulder, Colorado; and Sunnyvale, California are involved in a GFOA pilot project to develop interactive risk-aware forecasting coupled with budget models that indicate the risks of meeting individual budgetary line items. These models extend the experience of Colorado Springs by providing the ability to interact with the model and simulate the effects of uncertain revenue on the budget itself. Every time the forecast is changed, thousands of numbers are instantly run through the worksheet to estimate the chance of a shortfall in each budget line item.

Researchers have found that interactive, simulated experiences can be a highly effective communication tool for statistical information and for uncertainty and risk in particular.[2] Interactive simulation helps people better appreciate the role of uncertainty and chance in complex situations by allowing them to directly manipulate key variables and observe the results. It provides them with the awareness of uncertainty and the capacity to account for it, which leads to better decisions. Research has shown that simulation has led to better decision making for financial situations like planning individual retirement and perceiving the risk-return profile of investments.[3] It stands to reason that simulation could lead to better budget-related planning and decisions for public finance as well.

Monte Carlo Analysis. Simulation is based on an underlying statistical technique called Monte Carlo analysis, which is the mathematical equivalent of shaking a ladder before you climb it. Just as shaking a ladder bombards it with random physical forces to test its stability, Monte Carlo analysis bombards a mathematical model with random inputs to demonstrate what happens before decisions are carried out.[4] To illustrate, imagine you have a model where forecasted revenues minus planned expenditures equals a budget surplus or deficit. A Monte Carlo analysis could show the probabilities of ending up with surpluses or deficits by randomly varying the amount of revenue received, where the range of variation is based on historical experiences or even expert judgment. If your planned expenditures show an unacceptably high probability of a deficit, that might lead to different revenue or expenditure policy decisions.

The SIPmath Model's Reception in Redmond

The City of Redmond has only been using the model for short time, but finance staff has been able to come to the following preliminary judgments about the impact of risk-aware budgeting:

- **The model raises useful questions about the budget and forecasts.** Oftentimes, probabilistic thinking and analysis do not give a direct answer. Instead, they raise important questions that might not have come up otherwise. In Redmond, the model has raised awareness of the risks inherent in any forecast and helped the staff think about the relative risks that different revenue sources pose.
- **Specifying probabilities makes risk more tangible.** The model describes risk in quantitative probabilistic terms. This allows risk to be communicated visually through data graphics and to be discussed in precise terms (e.g., it is 80 percent likely) rather than vaguely (e.g., it will probably happen). This has helped the staff more specifically illustrate the risks of forecast options when discussing revenue estimates with executive management and the city council.
- **Risk-awareness improves forecasting ability.** Research shows that forecasters who describe forecast risk in precise probabilistic terms are more accurate than those who don't.* In Redmond, the model has prompted the city to reconsider the degree of conservatism in its forecast.
- **Is risk-aware forecasting and budgeting scalable?** Thus far, Redmond's project has been focused on its water utility. Redmond's great unanswered question is the difficulty of expanding risk-aware forecasting and budgeting to include other variable and revenue sources. With traditional Monte Carlo simulation, Redmond would have needed to build a single large model to cover all its departments. One of the biggest advantages of the SIPmath standard is that it allows departments to be modeled separately, whereupon the results may be rolled up to an overall budget. The city is looking forward to experimenting with this next level.

For Redmond's pilot, the next step is to better integrate risk-aware forecasting and budgeting information into presentations with elected decision makers in late 2016.

** According to an extensive study called "The Good Judgment Project." See: Philip E. Tetlocl and Dan Gardner,* Superforecasting: The Art and Science of Prediction *(Crown Books, 2015).*

Thanks to the increased power of personal computing devices and the new open-source SIPmath data standard, which runs in Excel without the use of add-ins or macros,[5] anyone with access to Microsoft Excel can now perform Monte Carlo analysis. The standard represents uncertainties as columns of thousands of possible outcomes in Excel or other environments. Forecasts expressed in this way can be made use of in the sense that unlike a bell-shaped picture, the columns of numbers may be run through a city budget spreadsheet to determine the chances of meeting specific budget line items. Furthermore, the forecasts are auditable, as they are made up of numerical data.

> Conservative and best-estimate forecasting strategies put the finance officer in a lose-lose situation.

This article will show the interactive SIPmath simulation model being developed by the City of Redmond, Washington (with numbers that are for illustration only). You may wish to download the model from www.gfoa.org/forecastbook now so you can follow along with the discussion. The first section of the model is the historical forecast and actual revenue data. (See Exhibit 2.) The historical differences between forecast and actual revenue are used to estimate potential future variations. The model also highlights when the city's forecast has been too optimistic or too conservative. Most forecast methods will have some bias built in. When assessing the risk inherent in the

Exhibit 2: Historical Data in Redmond's Model

Historical Data (in millions)			
2016	**Forecast**	$5.75	
	Conservative		
	Optimistic		
	Forecast	Actual	
2006	3.32	3.50	5%
2007	3.72	4.10	10%
2008	4.60	5.50	20%
2009	5.15	4.82	-6%
2010	5.10	5.20	2%
2011	5.23	4.80	-8%
2012	4.10	4.00	-2%
2013	3.96	4.40	11%
2014	5.50	5.90	7%
2015	5.25	6.50	24%

forecast, recognize which direction the bias goes so it can be accounted for. In Exhibit 2, the forecast model produces a conservative forecast in most years. The percentage difference between the forecast and actual, shown in the rightmost column, demonstrates that bias is often quite pronounced.

The next part of the model corrects for the bias. It is important to have an unbiased estimate as the starting point for assessing risk. For example, if the forecast already has a strong conservative bias, we might overestimate the risk of a revenue shortfall. Redmond's model corrects for this risk using a technique called regression analysis. Exhibit 3 plots the forecasted revenue against the actual revenue. If the forecast were completely accurate, all of the historical data points (the red diamonds) would fall on the blue dotted line, because that is where the forecast and the actual are equal. To be unbiased, the forecast does not have to be completely accurate, but the data points must be equally distributed on both sides of the line. As we see, three of them are on one side of the line and seven are on the other. Accordingly, the model estimates the amount of bias that is likely in Redmond's current forecast. The solid black line (the regression line) goes through the middle of all of the data points. The forecast (the yellow dot) is placed on the line to show the best estimate of the next period's revenue.

Exhibit 4 shows where the forecast bias is corrected and where we take the first step toward estimating risk. We can see that Redmond's new forecast is $5.75 million, as shown in the "deterministic" cell on the upper right. When we simulate an unbiased forecast, the best-estimate forecast becomes $6.03 million, also in the upper right. Also, based on the historical level of accuracy, the model then automatically calculates a prediction interval. In Exhibit 4, the level of confidence for the prediction interval has been set at 90 percent, and the bounds of the interval are $5.12 million and $6.93 million. This means that there is a 90 percent chance that actual revenue will be between these two numbers. The interval and the dial indicator on the left change as the user changes the level of confidence they wish to operate at, and as they change their forecast.

The final section of the model shows the implications of revenue volatility for budgeting. First, the model asks the user to divide the budget into prioritized tranches of expenditures. For example, making a scheduled bond payment might be a very high priority. Buying a new asset that is not strictly necessary for the continued viability of public services (e.g., new

Exhibit 3: Recognizing and Correcting for Bias in the Forecast

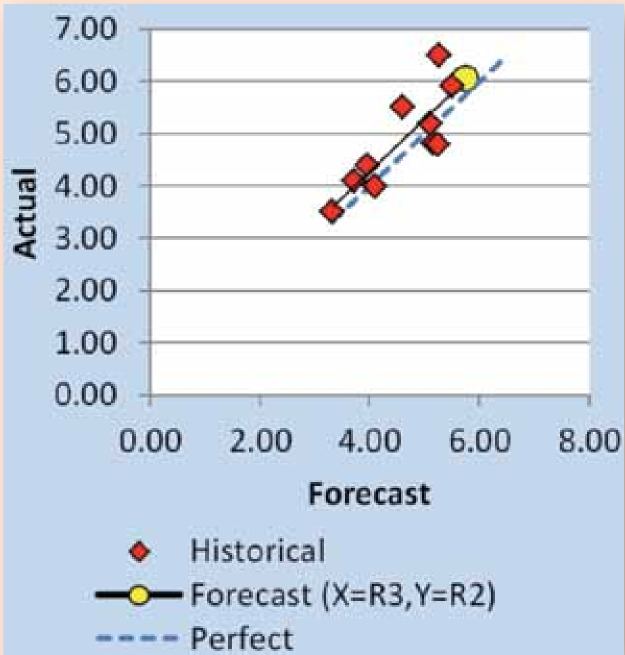

Exhibit 4: Simulated Unbiased Forecast and Prediction Interval

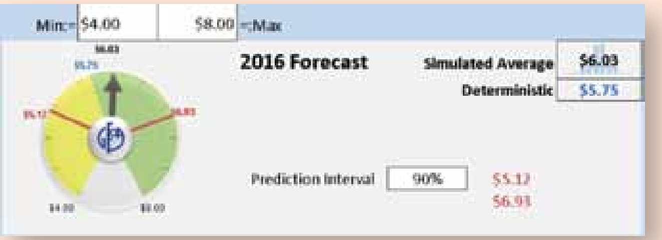

Exhibit 5: Point Estimate Revenue Forecast and Expenditure Budget

○ Risk-Aware ◉ Point Estimate

Point Estimate Revenue Forecast		$5.75
$5.75	Total Budgeted Expenditures (sum of items below)	Shortfall
$0.40	Debt Payments (Principal and Interest)	No
$2.49	Operating Expenses - Purchased Water	No
$1.96	Operating Expenses - General Operation	No
$0.90	Operating Expenses - Depreciation	No

Exhibit 6: Risk-Aware Forecast and Expenditure Budget

◉ Risk-Aware ○ Point Estimate

Risk-Aware Revenue Forecast		$6.03
$5.75	Total Budgeted Expenditures (sum of items below)	Chance of Shortfall
$0.40	Debt Payments (Principal and Interest)	0%
$2.49	Operating Expenses - Purchased Water	0%
$1.96	Operating Expenses - General Operation	2%
$0.90	Operating Expenses - Depreciation	31%

carpeting for offices) might be a relatively low priority. Then, the user has the option of seeing how this spending compares to the traditional point estimate forecast, where the risk posed by volatility is not taken into account. In Exhibit 5, we see how Redmond has divided its water utility services into four tranches, where making debt payments is the highest priority and covering depreciation is the lowest priority. Because the total expenditures are $5.75 million and the revenue forecast is $5.75 million, there is no shortfall expected. In the traditional mode of budgeting, this would be the end of the story.

Want to Learn More?

If the experience of Colorado Springs, Redmond, and the other pilot cities has you curious to learn more, here are some options:

1) Develop your own bell curve, like Colorado Springs did. The procedures for developing the curve are not difficult. Go to www.gfoa.org/forecastbook for step-by-step instructions and a spreadsheet.

2) Download a demonstration of the Redmond model at www.gfoa.org/forecastbook. Try changing some of the parameters to observe the results in order to get a better understanding of how risk-aware forecasting and budgeting work.

3) Stay tuned to GFOA for more results from the GFOA pilot project on risk-aware budgeting and forecasting. Assuming the cities continue to get positive results, we will provide more information on how to develop risk-aware forecasting and budgeting practices, including templates.

With the risk-aware budget, the model uses Monte Carlo analysis to simulate the likelihood of realizing a shortfall in each tranche. Factoring in the historical volatility of the city's revenue, we can see the chances that the city won't be able to cover its planned expenditures, line-by-line, and a spending plan that looked solid before has its vulnerabilities revealed. Exhibit 6 shows that there is a distinct possibility that forecasted revenue will not be sufficient to cover the costs of depreciation, even if we base the analysis on the higher, unbiased, best-estimate forecast. This insight might lead the city's decision makers to choose less ambitious spending plans or to consider alternatives to make sure they can fully fund the upkeep of their infrastructure. Since the SIPmath model is fully interactive, increasing Debt Payments to $0.50 would cause the model to instantly run 1,000 trials to reveal increased chances of shortfall in General Operation and Depreciation of 3 percent and 37 percent, respectively.

CONCLUSIONS

Explicitly considering risk in both forecasting and budgeting allows finance staff and budget decision makers to get the benefits of both conservative and best-estimate forecasting. The risk-aware approach allows the finance officer to share his or her best estimates of how the future will unfold, but it also encourages decision makers to adopt a budget that still balances if the future unfolds differently than expected. New technologies and standards allow risk-aware forecasting and budgeting to be performed in an everyday spreadsheet without a great deal of statistical knowledge. ❚

Notes

1. Many phenomena take this shape, including the heights of American men, standardized test scores, and even how close people park to the entrance at shopping malls. Example of normally distributed phenomena taken from: Charles Wheelan, *Naked Statistics: Stripping the Dread from Data* (W.W. Norton & Company, 2013).
2. Robin M. Hogarth and Emre Soyer, "Using Simulated Experience to Make Sense of Big Data," *Special Collection of MIT Sloan Management Review: Making Better Decisions*," Winter 2015.
3. Ibid.
4. Sam L. Savage, *The Flaw of Averages: Why We Underestimate Risk in the Face of Uncertainty* (Wiley, 2012).
5. SIP stands for stochastic information packet. "Stochastic" is another term for Monte Carlo analysis. SIP was developed by the non-profit firm, ProbabilityManagement.org.

SHAYNE KAVANAGH is senior manager of research at the Government Finance Officers Association in Chicago. He can be reached at skavanagh@gfoa.org. **SAM L. SAVAGE** is executive director of ProbabilityManagement.org; author of *The Flaw of Averages: Why We Underestimate Risk in the Face of Uncertainty*; consulting professor at Stanford University; and a fellow of the Judge Business School at the University of Cambridge. He can be reached at savage@stanford.edu.

If you want to learn more about how to apply these techniques and develop your own models, Sam Savage and ProbabilityManagement.org will be offering seminars that will provide an introduction to probabilistic modeling in Palo Alto, California (April 11, 2016), Washington DC (June 8, 2016), and Chicago (June 9, 2016). Government organizations can attend for a discounted rate. Visit probabilitymanagement.org for details.