

**Mukul Pareek, CISA, ACA, AICWA, PRM**, is a risk professional based in New York, USA. He has more than 20 years of audit and risk experience in industry and financial services. He is copublisher of the Index of Cyber Security, [www.CyberSecurityIndex.org](http://www.CyberSecurityIndex.org), and is looking for practitioner comments, feedback and opinions on quantitative measurement of technology risk. He can be reached at [mp@pareek.org](mailto:mp@pareek.org).

## Technology Risk Measurement and Reporting

It is difficult today to think of operational risk without thinking about technology risk. Workflow-based applications, system-driven notifications and databases with web front ends are proliferating, and they make operational processes indistinguishable from the systems on which they run. A failure in the process of creating a loss event can almost inevitably be tracked down to a technology control that was not designed well or that failed to operate. For instance, high-profile incidents at Barings<sup>1</sup> and SocGen,<sup>2</sup> among other places, showed inappropriate access to systems as a significant causal factor.

This expansion of the risks emanating from technology has significant implications for the technology risk manager, who is not only answerable for the traditional responsibilities of information security and data protection, but is also (welcomingly) involved in a multitude of business process discussions relating to access controls, segregation of duties, approval hierarchies, notifications, and automated communications to clients, vendors and insiders. As a result, technology risk is considered a large enough source of risk to often merit separate departments and budgets that may exceed the budget of the core operational risk function itself.

Yet, when compared to market and credit risks, the estimation, measurement and reporting of technology risks remains an undeveloped discipline. Risk measurement tools available today to the technology risk manager are at best not much more than crude directional indicators of risk. The measurement and communication of technology risk continues to remain an art, and is far from evolving to a science. The available tools—risk and control matrices with varying levels of risk and control granularity; red, amber and green dashboards; heat maps; quadrants; and other similar nonquantitative measurements of risk—do not come close to the sophistication of the tools available to market and credit risk professionals.

This article attempts to identify better ways of communicating risk by drawing parallels from

the more advanced disciplines of market and credit risk. Therefore, technology risk is looked at, within this article, as a significant subset of operational risk. This article revisits how risk measurement and quantification currently work for market and credit risk, and looks briefly at operational risk modeling as it is used for compliance with the Basel framework.

### CONTRASTING MARKET AND CREDIT RISKS FROM TECHNOLOGY RISKS

It is important to recognize the differences between financial risk and technology risk. This is necessary, as it is because of these fundamental differences that technology risk measurement continues to defy objective measurement when compared to financial risk. The key differences are:

- **Risk premiums**—Investors are paid positive risk premiums for taking on market and credit risk. For technology risk, there are no rewards for taking on risk except a possible avoidance of some unknown downside (and possibly avoiding the cost of the control). A fund manager may make a risky investment and earn a return greater than the benchmark index, and this is easily understood by management and the media alike. In contrast, it is difficult for IT risk managers to explain that they spent US \$2 million on information security and have no credible means to explain either the risk or the reward.
- **Relative importance**—In the financial services sector, market and credit risk dominate. These types of risk can make an institution go out of business. Operational risk or systems risk events may negatively affect an enterprise, but are unlikely to make one close its doors, except in the most extreme cases.
- **Availability of hedges**—Most market and credit risk can be offset by acquiring positions in other securities. There are no easy hedges for technology risk, other than implementing internal controls (though some insurance protection can now be purchased for a limited set of scenarios).



**Do you have something to say about this article?**

Visit the *Journal* pages of the ISACA web site ([www.isaca.org/journal](http://www.isaca.org/journal)), find the article, and choose the Comments tab to share your thoughts.

Go directly to the article:



## Enjoying this article?

- Read *IT Control Objectives for Basel II*.

**[www.isaca.org/research](http://www.isaca.org/research)**

- Consider Risk IT.

**[www.isaca.org/riskit](http://www.isaca.org/riskit)**

- Learn more about risk management and risk assessment in the Knowledge Center.

**[www.isaca.org/knowledgecenter](http://www.isaca.org/knowledgecenter)**

- **Measurement**—Market and credit risk can be measured and reported using value at risk (VaR),<sup>3</sup> exposures, limits and other quantitative tools. But, technology risk is difficult to measure. Most risk managers find it difficult to get beyond subjective red, amber and green indicators and their equivalents.
- **Fungibility**—In the financial markets, assets are identical and carry the same risk that can be hedged. On the other hand, if a company wishes to hedge against a particular technology asset (e.g., routers) being compromised, it is difficult to do. This is because, on average, only some of the total global population of that asset will be compromised, and not all of them at the same time. The contrast with financial risk is that when, for example, a currency goes down, it affects all investors holding the currency, not just a few. In the technology risk domain, the realized risk for a firm is binary, i.e., whether an adverse event happens or not.

Despite the previously mentioned differences, operational and technology risk contain many of the same elements as market and credit risk. The next section discusses risk measurement and reporting in the market, credit and operational risk worlds with a view to understanding the common elements and to attempting to identify learning opportunities for measuring and reporting technology risks.

### PARALLELS WITH MARKET, CREDIT AND OPERATIONAL RISKS

Financial risk managers look at risk in three broad categories: market, credit and operational risks. Risk calculations under each of these categories drive regulatory capital and economic

capital, and help firms manage their capital to their risk appetite and, as a corollary, manage capital levels to a desired credit rating from rating agencies.

### Market Risk

Market risk is the risk of losses arising from movements in market prices, including the risk of loss from changes in prices of financial instruments, foreign exchange rates and commodities. What one needs to calculate market risk is essentially position and price data. Using these, one can calculate correlations, volatility and the oft-quoted VaR number. At first glance, technology risk may appear to be so fundamentally different from market risk that any parallels may be difficult to draw.

A key difference is confidence levels. In the world of market risk, generally the question is framed differently in terms of confidence intervals: What is the worst loss with, say, a 95-percent level of confidence? To answer this question, one must consider the estimated future probability distribution of all possible outcomes and look at the bottom fifth percentile outcome. On the other hand, the question often asked of the technology risk manager is about the worst that can happen. Perhaps when talking of technology risk, one should consider thinking in terms of confidence levels. That is, instead of focusing on the worst case, which skews conversations to a scenario considered improbable by business managers, it is important to talk in terms of plausible scenarios with a certain level of confidence.

For example, if it is determined that the probability of losing one or more laptops during the year is 1 percent, one can say at a 95-percent level of confidence that we will not lose any laptops. At a 99-percent level of confidence, though, one or more laptops are likely to be lost. If an organization's senior management wishes to operate at a 99-percent level of confidence for avoiding a risk, it may choose to mitigate this risk by encrypting laptops. The level of confidence essentially reflects the management's risk appetite, or the level of risk with which it is comfortable.

### Credit Risk

Credit risk is the risk of loss in asset values from the downgrade or default of parties who owe money to the organization. These are generally loans, bonds and counterparty exposures for derivative positions. Interestingly,

while there are many models to measure credit risk at the portfolio level, they are all ultimately also distributional approaches in which an expected loss distribution is calculated and risk is measured in terms of confidence levels at a given time horizon (usually one year). Expected losses are the product of exposure at default (EAD), i.e., the amount owed by a counterparty; the probabilities of default (PD); and the loss given default (LGD), which accounts for partial recoveries. Expected losses are a product of these three variables, i.e., expected losses are equal to  $EAD \times PD \times LGD$ .

In the technology risk world, the PD corresponds to the probability of the risk materializing, and the LGD implies the actual loss caused if the event occurs. Exposures for technology risks can be measured in a slightly different and creative way:

1. **Exposure**—In the world of market and credit risk, exposure is measured by the size of the positions, the size of the portfolio or the monetary amount of the exposure. Technology risk is not measurable in the same way, but it is not completely immeasurable either. Technology risk practitioners are generally averse to making assumptions, and this reluctance often stems from the fact that they are more comfortable with precision.

This is where a shift of mindset may help—made easier by looking to the developed disciplines of market and credit risk. Innumerable assumptions and models underlie the measurement of these risks. These assumptions are acceptable to management, as well as to regulators and central bankers. For example, illiquid securities may have to be valued using assumption-based models. Credit risk exposure of a derivative contract can be estimated only using a distribution of the values it may take in the future, and these estimations rely upon a large number of assumptions. Similarly, the probability of the default of issuers is connected to rating classes, ignoring the unique financial characteristics of each issuer. Recoveries given default are equally assumption-driven estimates, given that they have varied extensively based on the business cycle and the industry. In other words, estimations and assumptions may provide an acceptable basis for measuring risk so long as there is a conceptual basis for the same.

Extending the analogy to technology risk, one must think a bit differently about the kind of exposures that technology risk managers face. A broad categorization includes the following largely comprehensive list of technology risk:

- *Information leakage* causing reputational and monetary losses
- *Business continuity risk* from failed systems and processes
- *External attacks* on infrastructure that may lead to either disruption or data losses
- *Process and workflow deficiencies* that allow fraud, theft or data leakage

The last category is particularly large in scope, as it includes a wide range of possibilities. An example is a poorly designed system that allows losses or fraud to occur undetected as the underlying system's workflows are not designed with the appropriate controls in place. This includes things such as absence of segregation of duties in business processes and risk created from missing IT general controls.

For each of the exposures, the next step is to determine what a firm's exposure is to that risk factor. For market and credit risk, exposures are measured in monetary terms or by the value of the parameters (called betas), indicating the risk sensitivities of the portfolio to underlying risk factors. The challenge with technology risk exposures is that an industrywide standard means of measuring and reporting such exposures might not exist. Nonetheless, at the firm level, this does not need to be a complex effort. In fact, it could be a positive exercise, as the firm can decide in which units to measure the risk exposure.

A measure of exposure could include the number of sensitive data records, the number of critical applications exposed to the Internet, the number of corporate e-mail accounts and the number of servers hosting critical applications.

Therefore, exposure for technology risks should be measured in terms of the drivers of such risk, possibly expressed in terms of numbers or whatever best expresses the extent of the risk. Market risk and credit risk have the great advantage of the risk analyst being able to measure all exposures in monetary terms. With the exposures for technology risk, it may not be desirable to always do so, as it may hide the true nature of the risk.

**2. The error rate of the control (or probability of default in the credit risk world)**—The extent of the exposure is effectively the level of risk facing the organization if there were no controls. This exposure needs to be offset by controls that are well designed and effectively operated. In the credit risk world, exposure is offset by mitigating factors such as collateral or guarantees. In the same way, for technology risk professionals, the exposure is offset by the existence of effective controls. For example, a firm may have a large number of servers that connect directly to the Internet, creating an exposure to the risk of Internet-based attacks. The presence of the right intrusion detection system (IDS) and other controls may effectively negate the risk, in which case the residual risk, or the net exposure, after taking into account the controls in place, may be zero. A related question that arises correlates with the measurement of the effectiveness of controls, i.e., how does one convert effectiveness into an error rate? Again, the time-tested techniques of risk-based audit and relying on sample testing provide a plausible answer. A randomly selected sample reveals the expected error rate of the entire population (the accuracy of the estimate is a function of the sample size, which can be drawn based on the level of confidence desired in the error rate). If one knows the error rate of a particular control, one can estimate the true population parameter. For example, if the control is expected not to prevent or detect the targeted process failure one time in every 50 transaction instances passing through the control, the true population parameter is 2 percent. At this point, the exposure and the error rate of the performance of the control are known; therefore, the actual number of occurrences of the control failure during a given period can be determined.

**3. Loss given control failures (akin to loss given default in credit risk)**—One could potentially stop at the previous point. But, to take things one step further, one can determine the loss given control failure. This needs to take into account the fact that not all control failures will lead to losses.

The following example illustrates what has been discussed in this section. Assume an organization has an exposure to losing confidential information through the corporate e-mail system, and the organization's rule-based e-mail surveillance can successfully block 90 percent of such outgoing e-mails. The organization's exposure is the potential number of e-mails that contain sensitive information, but are not blocked by the

e-mail system. If any such e-mail that is now outside of the organization's control gets compromised, the organization loses, say, US \$100,000. But, not all such outgoing e-mails are maliciously sent out (i.e., most will not be compromised even if they leave the firm's protected perimeter), and one may estimate 0.01 percent of the outgoing e-mails to have been sent out with malicious intent. One can now calculate the expected loss (Gross Number of E-mails Leaked  $\times$  0.01 percent  $\times$  US \$100,000) and the actual loss at different confidence levels, which would be akin to VaR (e.g., using a distribution based on the Poisson distribution, which needs just one parameter—the mean that was just calculated). Based on this information, a case can be made to management to increase the effectiveness of the e-mail surveillance system (by increasing its sensitivity, which will, in turn, increase the false positives and will take time and, therefore, money to resolve) to a level at which management is comfortable with the remainder of the risk.

### Operational Risk


Operational risk is the risk of loss resulting from inadequate or failed internal processes, people and systems, or from external events. (In this article, reputational risk is considered to be a part of operational risk, even though the Basel framework that applies to financial institutions specifically excludes it.) As argued earlier, technology risk is a key component of operational risk.

One criticism of operational risk is that it is a top-down approach largely regarded as disconnected from the realities of everyday technology risk management. It is focused on calculating capital adequacy numbers for regulators and management, and does not help identify real actions that can be taken to address risk. Yet, an extremely important tool that operational risk practitioners employ is scenario analysis that produces estimates of loss frequencies and impact, and that, at the same time, increases the awareness of risk among the managers who participate in the scenario analysis exercises.

The Basel framework requires a generic implementation of the advanced measurement approach (AMA). A generic model for modeling operational risk (**figure 1**) works as follows:

- Operational risks are understood as the product of the frequency and severity of loss events:
  - Frequency is the number of loss events happening during a time period.
  - Severity is the realized loss impact of the event.

**Figure 1—Modeling Operational Risk Losses**



- Modeling frequency; the frequency of losses is modeled using an appropriate distribution (usually the binomial or Poisson). These distributions need only a couple (or even a single) parameter(s), which can be estimated as follows:
    - Expected Loss Frequency = Loss Probability × Number of Events, Transactions, etc.
    - Example: Assume that the probability for credit card fraud equals 0.01 percent of all credit card transactions and that the number of transactions in the loss horizon is 1,000,000. As such, the expected loss frequency, or  $\lambda$ , is equal to 0.01 percent × 1,000,000 = 100.
    - With just a few assumptions, one has a frequency distribution.
  - Modeling severity; the severity of losses is estimated using a lognormal distribution with a mean and variance ( $\mu$  and  $\sigma$ ) estimated using focus groups, scenario discussions, etc.
  - The product of the two is obtained by using Monte Carlo simulation, i.e., pick one random value from the frequency distribution and one from the severity distribution, multiply them, and the result is a data point. Repeat this exercise multiple times to generate enough data points to build a loss distribution.
  - The loss is calculated at the fifth or the first percentile, depending upon the confidence level desired.
2. Do not make the worst-case scenario the focus of risk communication; note the worst case, but do not make it the default case.
  3. Set the risk appetite. At what level of *ex ante* probability is the organization or management comfortable living with a certain type of risk? Plan for control implementation, allowing for the risk to be realized with this probability.
  4. Clarify and outline risk exposures and underlying drivers in numerical terms.
  5. Evaluate controls with an eye on error rates—empirically known or determined by sample testing. Track these over time.
  6. Use scenario analysis to not only discover or quantify risks, but also as a forum to educate managers on possible risks.
  7. Plug data gaps with judgment-based assumptions when measuring and reporting risk, but identify them clearly.
  8. Modify assumptions iteratively when observations turn out to be different.
  9. Try to build a loss distribution to determine technology risk losses at different confidence levels, and improve it over time.
  10. Add to the exposure universe by being observant of losses experienced by industry peers.
  11. Finally, when communicating risks, use jargon-free plain language, focus on reasonable accuracy as opposed to absolute precision, be forward-looking, and facilitate discussion and judgment.

### CONCLUSIONS FOR TECHNOLOGY RISK REPORTING

The above discussion highlights the tools used by risk managers in the financial world, and these contain a number of messages for the technology risk manager. The key takeaways here are:

1. Think in terms of confidence levels, and in terms of probabilities of risk realization.

While the evolution of technology risk measurement will happen over time, it is important for the technology risk manager to be aware of how risk measurement works in other risk disciplines, and possibly to borrow a tool or two to advance the science, while retaining the art.

## REFERENCES

Basel Committee on Banking Supervision, *Basel II: International Convergence of Capital Measurement and Capital Standards: A Revised Framework—Comprehensive Version*

Hull, John C.; *Options, Futures and Other Derivatives*, Prentice Hall, 2005

Index of Cyber Security, [www.CyberSecurityIndex.org](http://www.CyberSecurityIndex.org)

## ENDNOTES

<sup>1</sup> Bank of England, “Report of the Board of Banking Supervision Inquiry Into the Circumstances of the Collapse of Barings,” 18 July 1995. The report identified various causes for the collapse of Barings. A key theme was the lack of segregation of Nick Leeson’s duties; he was able to enter and approve transactions and reconciliations without supervision.

<sup>2</sup> Societe Generale, Summary of the Court’s order, [www.societegenerale.com/sites/default/files/documents/Summary\\_of\\_the\\_committal\\_order.pdf](http://www.societegenerale.com/sites/default/files/documents/Summary_of_the_committal_order.pdf). In January 2008, Societe Generale, France’s second-largest bank, announced that Jerome Kerviel, a trader on the futures desk, had lost US \$7.1 billion of the bank’s money in rogue trades. The fraud was concealed because Kerviel was not only responsible for entering into trades as a trader, but also had incompatible responsibilities to record the trades in the books. He entered fictitious deals to cover his tracks, only to subsequently cancel and even simply erase them in the bank’s computerized database—a case of inappropriate systems access.

<sup>3</sup> Value at risk is a common measurement of risk, particularly market risk. In the context of market risk, what it requires one to know is a future distribution of portfolio returns, or of portfolio value. A distribution of future returns or values is nothing but a listing of all possible future outcomes for a portfolio of assets at a certain time horizon, often two weeks. The VaR is simply the loss at the fifth percentile level. It answers the question: At a given level of confidence, and over a given period of time (and given assumptions about volatilities, correlation and distributions), what is an estimate of the loss over a fixed time horizon that would not be exceeded with that given level of confidence? A VaR number expressed at 99-percent confidence implies a 1 percent chance (which really means two or three days in a year) that this loss estimate will be exceeded. VaR does not answer the question of the worst-case loss, but is a distributional approach to measuring risk at a certain probability level. If one knows the distribution, everything about risk is known—at least in theory. Given that the distribution is often assumed to be normal, VaR just ends up being a multiple of standard deviation.