

Toward a Prime Metric: Operational Risk Measurement and Activity-Based Costing

by Michael Mainelli

In their quest to assign capital appropriately, banks need the ability to compare two products in a way that objectively ascertains which is operationally riskier and by how much. Perceptions of changes in these risks can be used to predict future risk. *Enter the prime metric to save the day... (if only it were that easy).*

Operational Risk Matters

Regulators believe that capital requirements for market, credit, and operational risk are a key regulatory mechanism for banks. "Capital is pivotal to everything that a bank does and changing it—and we believe Basel II could change it dramatically—has wide-ranging implications for bank management and bank investors."¹ A bank's capital requirements affect its risks in several ways, including:

- Constraining a key performance measure—return on equity.
- Influencing a bank's ability to lend and spend.
- Limiting dividends and capital repatriation.

Thus, by changing capital require-

ments, a bank's behavior toward risk should change as well.

For the first time, Basel II permits regulators to levy a capital charge for operational risk through one of three approaches:

1. Basic indicator approach—capital is calculated on gross income.
2. Standardized approach—there is a different indicator for each line of business, e.g., corporate finance, trading, retail banking, commercial banking, payment and settlement, retail brokerage, asset management, etc.
3. Internal measurement approach—banks calculate their expected loss by line of business, and regulators apply an additional factor.

Basel II has indicated that advanced risk-transfer techniques, especially insurance, will lessen capital requirements only for those banks that use advanced-measurement techniques within the internal measurement approach.

Operational risk does matter. In 2001, for instance, Schroder Salomon Smith Barney estimated that Deutsche Bank could free up to €5.8 billion from market and credit risk capital under Basel II, reducing its capital requirement by 27%, from €21.6 billion to €15.8 billion. However, at €2.8 billion, up from nil, operational risk would be nearly 18% of Deutsche Bank's remaining €15.8 billion capital requirement. For banks that can convince regulators of the appro-

© 2004 by RMA. Michael Mainelli, Ph.D., FCCA, is executive chairman, Z/Yen Limited, London, UK. Z/Yen helps organizations make better risk/reward management choices and undertakes strategy, finance, systems, marketing, and intelligence projects in a wide variety of fields, such as developing an award-winning risk/reward prediction engine, helping a global charity win a good governance award, and benchmarking transaction costs across global investment banks. The author would like to thank Giles Wright of Z/Yen for providing the background analysis.

priateness of their internal measurement approach, operational risk capital charges can be reduced and competitiveness increased.

While seeking appropriate internal measurement approaches, banks still lack a prime metric that gives the operational risk of a product or line of business. While there are a number of top-down, hierarchical checklist approaches or graduated risk scales (e.g., RAG—red, amber, green—reporting), no one seems to be proposing a comparative prime metric. Banks need the ability to compare two products in a way that objectively ascertains which is operationally riskier and by how much. Perceptions of changes to the inputs used to calculate the prime metric can be used to estimate future operational risk. Prime metrics exist for credit (risk of default) and market (value-at-risk) risk. Why not for operational risk?


The Prime Metric Problem


Basel II has introduced a definition for operational risk as “the risk of direct or indirect loss resulting from inadequate or failed internal processes, people and systems or from external events.” Operational risk has many pseudo-standard sub-taxonomies, such as *people* (e.g., workforce disruption, fraud), *process* (documentation risk, settlement failure), *systems* (failure, security) and *external risks* (suppliers, disasters, utilities failures). Categorizing risks is more empirical than analytical, and

DAY-TO-DAY OPERATIONAL RISK MANAGEMENT INVOLVES DECISIONS ABOUT OPENING TIMES, CLEANING STANDARDS, RODENT CONTROL IN DEALING ROOMS, SECURE ELECTRICITY SUPPLY, SECURITY CONTROLS, AND OTHER MANAGEMENT DECISIONS NOT SUITABLE FOR REAL-TIME SPREADSHEET ANALYSIS.

taxonomies of risk can seem arbitrary, overlapping, or in contradiction to each other. To an outsider, disasters such as Barings or AIB seem to be people or reputational

risks, while U.S. savings and loan fiascos or Argentina’s financial failures cannot be easily pigeonholed. Day-to-day operational risk management involves decisions





about opening times, cleaning standards, rodent control in dealing rooms, secure electricity supply, security controls, and other management decisions not suitable for real-time spreadsheet analysis. There is a tension between the top-down imposition of a capital charge and the bottom-up nature of these detailed decisions, as is evident from the following quote from a BBA document:

Risk information. The generation of complete and consistent risk data is a primary objective for many firms. A number have reported short-term gains achieved through the utilization and comparison of existing data, but an increasing number of firms are building specific operational risk identification and assessment systems. A good deal of the data is nonfinancial and subjective: e.g., key control standards, key risk indicators, the output from operational risk self-assessment. Firms also col-

lect internal operational risk loss and near-loss data.

Progress is slowed by the expense of implementation and the consequent requirement to demonstrate the business value of data collection in the near term.”²

Part of the problem with establishing operational risk metrics is that each set of metrics is designed to satisfy a different stakeholder view. There are at least three stakeholder viewpoints:

1. Regulators—looking at systemic risk and imposing capital requirements.
2. Debtors—being frequently influenced by ratings agencies and the security of their loans; also interested in risk-avoidance approaches, e.g., mutual operational risk insurance “clubs,” unsecured, subordinated debt issues,³ direct insurance provision, or catastrophe “opbonds.”
3. Shareholders, overlapping with bank management—wanting to see that things are

well run, typically by benchmarking against other banks; currently using operational risk benchmark clubs, operational risk methodologies, and professional services advisors.

Figure 1 attempts to summarize the available approaches in six categories, looking at each of the three views and distinguishing between the more frequent and less frequent risks. The regulator and debtor views do not give much guidance to an operational risk manager in a bank. The objective of most regulatory action is to minimize systemic risk. This focus preserves the banking species, but is not much help to the survival of any individual bank. Some of the long-touted mechanisms of the debtor view are starting to be taken seriously—for example, external, mutual, or internal insurances work if there is an informed market, a good negotiator, and a lot of experience, but they supplement, not replace, day-to-day operational risk management.

From the shareholder and management viewpoint, event-loss databases and benchmarking are beloved by consultants, and there are databases of operational losses as well as numerous “best practice” guidelines. However, these databases favor large, publicly known losses that are not of much value in daily management of operational risk. Internal risk management databases are biased by their small sample sizes, youth (as yet), and the lack of a management structure for operational risk in banks. RMA—The Risk Management Association has worked on the possibility of relat-

ing key risk indicators (KRIs) to losses using predictive risk/reward tools, such as Z/Yen's PropheZy. Yet even if the KRIs turn out to be predictive, a prime metric would still be useful in assessing the performance of one product line against another.

Activity-Based Costing and Industrial Operational Risk Management

Banks do not have a monopoly on operational risk. Risk managers are found in diverse non-banking environments, e.g., manufacturers, service firms, lawyers, accountants, hospitals (clinical outcomes), charities (care outcomes), or property companies. Risk managers fill diverse functional roles, such as production-line quality control, health and safety, insurance risk management, project management, business continuity or disaster planning, as well as general business operations. For example, an aircraft manufacturer runs serious operational risks, from designing through building to product service. An airline has operational risks from using the aircraft manufacturer's products, through sales to operations, freight, and security, to name but a few.

A crucial insight is that operational risk is intimately related to quality. Quality is frequently defined as "fit for purpose." There is a wealth of literature on quality management, ranging from inspirational "Rah-rah, total quality management is everything" through qualitative ISO9000 approaches, which bear more than a passing resemblance to banks' operational risk checklists. Some

A CRUCIAL INSIGHT IS THAT OPERATIONAL RISK IS INTIMATELY RELATED TO QUALITY. QUALITY IS FREQUENTLY DEFINED AS "FIT FOR PURPOSE."

banks already use semi-quantitative Six Sigma approaches that emphasize the importance of measurement. Despite the diversity of operational risk in industrial environments, one overriding principle in leading organizations is the importance of measuring variance in costs and quality.

High-risk processes tend to correlate with high-cost variances. In turn, high-risk processes tend to correlate with low-quality outputs. In quality-obsessive industries, constant measurement of cost variances is used to detect quality problems. Operating under the assumption that the financial systems must trap all process costs and relate them to outputs, industry uses these measures to manage operational risk. Activity-based costing systems are essential to quality measurement. Industrial companies attempt to include the full process cost, i.e., not just direct costs, such as raw materials, but also rework, wastage, scrap, disputes, legal fees, returns, environmental fines, and so forth. In many cases, industrial companies undertake *discrete* (most manufacturing environments) or *continuous* (most chemical and fluid environments) simulation modeling and then relate the models to costs. The activity-based costing systems then give managers costs per unit of output or costs per batch.

For a production manager looking at two assembly lines pro-

ducing the same product—say, cans of beans—the operational risk of each line can be assessed. Line A produces beans at an average cost per can of \$0.08. Line B produces beans at an average cost of \$0.10. Line A has a standard deviation of \$0.02 and Line B a standard deviation of \$0.01. While Line B does have a cost problem, it is the lower-risk line. Line A is the higher-risk line. Applying this approach to banking might have some merit. Let's begin testing the theory by looking at some data.

"In theory, there is no difference between theory and practice; in practice, there is."

—Chuck Reid

Z/Yen regularly benchmarks investment bank costs, head counts, and volumes to produce costs per trade covering:

- FX and money markets—global FX, currency options, money market.
- Equity and debt markets—European equities, SBL, bonds, repo, listed and OTC derivatives.
- U.S. securities transaction costs—equities, stock borrow loan, bonds, repo, options and futures.

Figure 2—four graphs of the cost-versus-volume curves for global foreign exchange, global money markets, European processed equities, and European processed bonds—shows that volumes have increased markedly

from 2000 to 2002, while operations costs per transaction are falling.

Figure 2 captures operations costs—trade capture, static data, presettlement, settlement, exceptions management, projects, management, and administration—in a standardized way against a set of agreed management accounting principles. They do not include middle-office or product-control costs or IT costs (though separate analyses are available). In three of the four markets—global money markets, European equities, and European bonds—“average” bank volumes in 2002 were volumes that only the largest processors handled in 2000. Each graph clearly shows significant cost decreases and a race to economies of scale.

So, if the manufacturing analogy holds true, then the

standard deviation of cost per trade should be lower for lower-risk products. The standard deviation of cost per trade needs to be normalized in order to permit comparison over products that have varying cost bases. One nor-

malization approach is to divide the standard deviation by the range between the minimum and maximum cost per trade; thus, a high normalized standard deviation over the range indicates a low-risk product. It is important to





Figure 3

Normalized Standard Deviation as a Valid Prime Metric

Annual Volume and Normalized Standard Deviation of Cost per Trade Between Organizations



state that the intra-bank prime metric described earlier ought to be applicable *within* a bank using a single bank’s internal activity-based costing system. Although interbank comparisons of normalized standard deviations might be an approach of use to a regulator, the comparison among banks shown in Figure 3 is meant to explore whether normalized standard deviation might be a valid prime metric. Figure 3 shows the normalized standard deviation of cost per trade across a “peer group” of banks.


The results are interesting. The method does seem to cluster perceived higher-risk products to the right and the lower-risk products to the left. Direct comparison of one product with its neighbor can be problematic—for example, are European cash equities riskier than global credit derivatives?

Some clearly low-operational-risk products, such as U.S. stock lending and U.S. bonds, seem to correlate nicely with the normalized standard deviation argument, and some perceived high-operational-risk products seem to correlate well also. On the other hand, some products do seem to be rated a bit riskier than experience might suggest, while others seem to be rated a bit less risky. But these normalized standard deviations are based on small numbers. The peer group bank numbers for each product range between six and 10 banks, depending on the products the banks traded in. Obviously, this small sample indicates that the approach toward a prime metric may be valid but is not conclusive. Further research on benchmarked interbank costs might be worthwhile, particularly if backed by regulatory interest. Even better, it

would be interesting to have research looking at intra-bank costs against a standard activity-based costing methodology.

The earlier graphs clearly show a strong relationship between volume and cost. Given the high fixed-cost nature of many banking products, volume fluctuations on their own can change cost per trade markedly through no fault of the operations. As higher volumes correlate with lower costs per trade, it might be interesting to explore a different variance, such as the variance from a cost/volume trend line. Assuming that the cost per trade fits a logarithmic volume curve, it is possible to examine the effects of volume using a “least squares fit” through points where $y = c \ln x + b$, c and b are constants, and \ln is the natural logarithm function. The key test now is how well the

Figure 4



banks' costs per trade fit the trend line—that is, r-squared (r^2). R^2 is the fraction of the total squared error that is explained by the volume. Thus values approaching 1 might imply that cost per trade is simply due to volume, or that the product had very low operational risk, or a bit of both. Values approaching zero might imply that operational risk is high. Figure 4 shows the same products ranked by the r^2 of the logarithmic trend line. Again, the results are interesting, and a few products have moved up considerably—for example, global IR derivatives, global credit derivatives, and global vanilla equity options. Meanwhile, a few products have moved down considerably—for example, global currency options, European bonds, and European cash equities. Depending on perceptions, Figure 4 might seem a closer fit to market perceptions of operational risk than Figure 3, although there is some broad agreement between both graphs on the implied risk in some product lines and some complete concurrence—for example, the least risky product line, U.S. stock lending, as well as European stock lending, European exchange-traded derivatives, and European

repo. On the other hand, U.S. repo has moved up and would generally be seen as a less risky product line. It is worth reiterating that these numbers need to be used with great caution and are based on a very small sample of banks. What seems worth further examination is the use of volume-based r^2 or normalized standard deviation as an operational risk prime metric.

Not Cost, but Swings in Costs

Despite the ubiquity of this cost-variance approach in industry, it is rare to find intra-bank product comparisons used to ascertain operational risk in banking. In our experience, a few banks have activity-based costing systems containing enough rigor to produce the intra-bank numbers, yet we have not seen that they examine the operational risk implications of variances in cost from period to period.

A good first step for a bank seeking an operational risk prime metric might be to develop “full” per-transaction costings and compare the variances by product line over time with loss profiles. As an aside, risk/reward option theory can be used to price both the cost of operational risk and the bene-

fits of options for change using similar skills to those used on the trading floor, such as Black-Scholes calculations. A good first step for regulators might be to commission detailed research on this approach. If the normalized standard deviation or r^2 for volume curves seems to hold up for larger samples, then some solid management accounting standards combined with semi-coercive benchmarking might go a long way toward providing an operational risk prime metric. Given a prime metric with associated standards, banks could sift the riskier from the less risky without having to wade through thickets of tick marks. □

Mainelli can be reached by e-mail at Michael_Mainelli@zyen.com.

References

- 1 "Time to Catch Up: Basel II—Modern Capital Rules for Modern Banks," Schroder Salomon Smith Barney, September 2001.
- 2 "1999 Survey—Operational Risk—An Industry Discussion Paper," British Bankers' Association, p. 4, www.bba.org.uk/bba/jsp/polopoly.jsp?d=146&a=491
- 3 "The New Basel Capital Accord: Making It Effective with Stronger Market Discipline," H.A. Benink and Clas Wihlborg, *European Financial Management*, Vol. 8, pp. 103-115, 2002.

Watch for more on operational risk coming in the July/August issue of The RMA Journal

“Proving the Value of Operational Risk Management,” by Patrick O’Neill, head of Operational Risk, America, BNP Paribas, looks at the business case for ORM.

“Getting Traction with KRIs: Moving Ahead,” by Jonathan Davies, director, RiskBusiness, updates an earlier *Journal* article, “Getting Traction with KRIs: Laying the Groundwork,” and gives a progress report on the global *KRI Framework Study*.