Op Risk Squared: Managing the Risk in Operational Risk Modeling

Mark Levonian

Office of the Comptroller of the Currency

OpRisk North America, New York, March 2012

I'm not likely to surprise anyone in this audience if I say that approaches to operational risk have changed a lot in recent years. Some of that change has been driven by Basel requirements, and developments like the Advanced Measurement Approach, or AMA. I know that my Japanese colleague Mr. Mitsutoshi Adachi, the Chair of the SIG Operational Risk subgroup, plans to talk more about Basel developments in his keynote remarks here tomorrow morning.

Basel and AMA have provided much of the impetus behind change – but not all of it. More fundamentally, it is the evolution of information technology and of the nature of financial services – what those services are, how they're managed, how they're delivered – that has really shifted the nature and sources of op risk. And with those changes, how we assess and manage op risk has changed too.

Op risk also has certainly become more prominent. Right now at the OCC, when we go through exercises to rank the various risks we see at the larger banks we supervise, operational risk has been number one, the top risk. Of course, these risk rankings depend on a number of things and are transient, and credit risk will always be big for banks. But the top position for operational risk is at the very least symbolic of the change in focus, by both banks and regulators.

At the same time, as operational risk has become more prominent, models have come to play a much more significant role in the assessment of operational risk. The reason is simple. In modern financial firms, models provide the tools for systematic analysis of data, in this case operational loss data.

Models can and should support operational risk management

In my view, Basel has done the industry a service by driving development of quantitative approaches to operational risk, and especially of pushing a systematic approach to operational loss data. Op risk models make significant demands on data, which has required new approaches to data collection in new areas. We've all watched and participated in the sea change in the op risk world on the data collection side, with the tremendous recent emphasis on collecting, organizing, and using more and better data.

But there is a potential down side of this too, with so much of the development in op risk modeling driven by the requirements of Basel and AMA. It can lead to a sense that the modeling is solely a compliance exercise, with little real value. I sometimes hear that from regulators, but more often I hear that tone from bankers. "Just figure out what will satisfy the regulators, and do it with as little work as possible, without interfering with the real business of the bank. Oh, and if possible, make sure it shows a fairly low capital requirement."

And that's a mistake. Operational risk should be approached like any other risk to be managed. As with other types of risk, models provide invaluable support to decision making.

Actually, that statement itself – that op risk is like other risks to be managed – is not accepted by everyone, so let me digress for a moment. There is a view that operational risk is different, specifically because there is supposedly no risk-return tradeoff. I reject that view. Organizations constantly make choices about activities to mitigate or control operational risk, choosing the quantity and quality of people, backup systems, legal support, and so on, all of which require commitment of resources. A firm can choose to spend less on these things, and face higher operational risk; or spend more on mitigation, and reduce risk. Spending on the mechanisms that address operational risk – redundant systems, better management information, more effective governance, better legal support – increases expenses and lowers profits. A firm could choose to spend less, pass the savings to the bottom line, and run higher risk. Or choose to spend more to reduce risk, and reduce bottom-line profitability. This may not be a textbook risk-return tradeoff, but it is fundamentally the same thing: risk and return move in the same direction.

Every business faces this tradeoff, and needs to make responsible decisions about how much to commit to mitigating risk. In the modern financial world, these kinds of complex decisions shouldn't be made without first-rate quantitative information support provided through sound modeling. Operational risk models should be developed and used in ways that foster good decision making around the management of risk. They are not, or should not be, just a means to satisfy a requirement for a regulatory capital calculation. I believe the truth of this is becoming more widely recognized over time, but some firms are farther down this path than others, and most have some distance yet to go. In some ways progress has been handicapped by the regulatory origins of much op risk modeling; the "pure compliance" view is a stigma that must be overcome.

Models introduce an additional risk to be managed

Recognizing that op risk models should be part of key decision-making in risk management may up the ante for some firms. If this isn't something you're just doing for regulators – if it's about effectively managing the firm – you have to ask yourself: what if the models are wrong?

I'm sure many of you have heard the famous statement by the statistician George Box that "All models are wrong, but some are useful." That is probably the most succinct argument in favor of the use of models. But Box followed this with a sentence that is much less familiar; I quote: "However, the approximate nature of the model must always be borne in mind…" If the first sentence is the reason for using models, the second is the reason for model risk management. And that is what I want to turn to now.

I assert that the use of models is crucial to good management in financial firms today. No firm can address the high volume of available data in a systematic way without modeling. But sometimes even the best models don't work as planned, and sometimes that can be costly. Although we don't generally think about it this way, this "model risk" could be viewed as a type of operational risk in itself – a failure of systems and processes. Thus my title, Op Risk Squared – the risk that systems used to measure operational risk – in this case op risk models – may themselves fail presents a kind of operational risk to be acknowledged and addressed. The

leading firms do recognize this, and put in place effective model risk management policies and practices that help protect against unpleasant surprises due to the failure of models or modeling processes.

OCC supervisory guidance can help

In that regard, the OCC issued supervisory guidance last year, OCC Bulletin 2011-12, developed jointly with the Federal Reserve, that updated earlier, long-standing OCC guidance on model validation. My sense is that this guidance may have received less attention in the operational risk modeling world than it deserves, and less than in some of the other risk disciplines. With that in mind, let me take a few minutes to describe some of the key elements of the guidance, and share some thoughts on how it applies to op risk modeling.

The guidance reflects what might for some be a new perspective, albeit one that has gained greater influence in the wake of the financial crisis. Prior supervisory guidance on models focused largely on validation. But the more modern view is that validation, while still of key importance, is only one part of a broader approach to managing model risk.

Actually, just characterizing model risk as a risk, similar to other risks, helps simplify and clarify thinking a lot, because it sets in on a familiar foundation, and keeps us from having to develop everything from a zero base. Because to manage risk, any risk, there are certain things we know have to be done.

First, identify the sources of risk, for example through an inventory of models, with appropriate documentation. Second, assess the magnitude of the risks, perhaps distinguishing between likelihood and impact; this might involve sensitivity analysis, to look at the potential impact of model failure on P&L. Third, take steps to control or mitigate the risks. The greater the potential impact, the more material the risk – the more rigor is necessary and expected in risk control activities. Just like with other risks.

The guidance emphasizes the importance of "effective challenge," of having somebody whose job it is to poke at the model and see if it breaks, who is supposed to be able to stand up and say "this doesn't work right".

Where does that effective challenge come from? Use is one potential source of challenge, because generally model users want models that work. That's why you see "use tests" figuring prominently in supervisory guidance like the AMA guidance produced by SIGOR and others. Another source of effective challenge can be internal audit, or model review. But validation remains a central and necessary source of effective challenge for most models.

Our supervisory guidance notes that for validation or any other form of challenge to be effective, there are three keys: incentives, competence, and influence. Where challenge is ineffective, generally one of these three is missing.

Proper incentives can come through independence, with independent validation or review. But that's a "can," not a "must." Other ways to create proper incentives are through the structure of compensation, through performance standards, and through corporate culture. Don't underestimate that last one; it's hard for good controls to overcome bad culture.

Beyond incentives, effective challenge requires competence – without adequate technical knowledge and skill, even challengers with the best incentives will be ineffective.

But clearly competence, while a necessary condition, is not a sufficient condition. Even competent challengers with the right incentives to provide vigorous challenge to models won't be able to get results if they lack that third element – influence. Influence means that if problems are found, action gets taken, problems are not ignored. Influence comes in part through explicit authority, internal stature, strong higher-level commitment and support, or most likely some combination of all of these.

Notice that I didn't say effective challenge requires creation of a distinct validation group. It can, but we often find that model developers are best positioned to do much of the validation work. They tend to be high on competence and influence (at least on models). But they are weaker on incentives in some cases, so where developers are doing validation work the

firm needs some secondary, supplemental source of challenge, such as independent review of their validation work.

Much of our supervisory guidance, both for model risk and for operational risk as a whole, emphasizes governance and controls. We're talking about managing risk, and we all know that works best when there are clear and robust internal processes. Validation and model risk management are ongoing processes, not once and done.

I'm fairly certain that my OCC colleague Carolyn DuChene will touch on governance issues in the next session this morning, so I won't go into details here. But I will note that an important element of effective model risk management, emphasized in the new supervisory guidance, is the assignment of clear roles and responsibilities around model risk management. These roles in include model ownership, risk control, and compliance. Someone must "own" each model, with clear accountability for use and performance. The risk control role covers many of the things I've just been talking about: assessment of model risk, provision of effective challenge, and the authority to restrict or limit model use and application as necessary. Compliance ensures that policies and procedures are effective and are followed.

Current areas of focus in op risk modeling and model risk management

As op risk models move beyond regulatory capital into effective risk management, I predict that model risk management will take on added importance. As this occurs, the supervisory guidance I've just been describing will be a valuable guide. Let me now briefly touch on just a few representative issues from our supervisory work in the area of operational risk modeling and model risk management. I am confident that many of these issues will get more attention and discussion as this conference proceeds.

Several key elements of model risk in the op risk space relate to data, and data integrity. This is a deep topic, and picking out only one or two examples runs the risk of suggesting that these are most important, or that others are less important. But think about the use of external loss data, as required under AMA. This generally requires mapping loss categories from the

external source data to corresponding loss categories in a firm's own data. Mapping is an art, with expert judgment required. Data also typically must be rescaled to be relevant.

Effective model risk management should then be asking questions like: What decisions were made in the mapping and scaling processes? On what basis? What alternatives were considered? What's the impact of the choices made on the resulting output of the operational risk model? Is there a process in place to periodically review these decisions? Who is responsible for that review process?

Or think about the common problem of setting a severity threshold for loss data capture, which is common at many firms. Was the threshold chosen through some careful systematic process that considered the impact on modeled outcomes? Or was the decision arbitrary, or based on convenience? Asking and answering these questions would be part of effective challenge.

Turning from data to a typical issue related to the structure of the op risk modeling framework, an obvious key decision is the selection of units of measure, such as the set of combinations of business lines and event types into which losses are categorized. Model risk management questions here could include: Does the framework have too few of these to create reasonably homogeneous units for which it makes sense to model severity and frequency? Or too many, so that the resulting loss data are too thinly spread to produce robust estimates of loss distributions within each unit of measure? A good model risk management framework should be looking at those structural decisions and their impact, how the decisions were made, and whether there is a process for periodically reassessing the selected structure.

An element that has received considerable attention in op risk modeling is the choice of dependency structure across units of measure, such as the choice among alternative copulas. Modeling results can be quite sensitive to the choice. Good model risk management should focus on the developmental evidence for key modeling decisions like this. Evidence that the resulting estimates are reasonably robust to alternatives provides valuable comfort, with common sense – for example, about which units of measure should be most closely related to one another – a valuable guide, and conservatism a backstop.

As in other types of modeling, validation has an important role to play. While outcomes analysis like backtesting should be a component of validation, recognize that most common copulas look similar in the center of the distribution, where the weight of the backtesting results will be. Outcomes will be thinner in the tails, which unfortunately is where alternative copulas can really be distinguished. So backtesting may be a weaker tool in this case. A viable alternative may be benchmarking, comparing the results of multiple models or approaches. Good, credible benchmarks are likely to be especially valuable in op risk modeling.

Stress testing and operational risk models

I can't conclude without a few comments about stress testing. Clearly this is a big topic in risk management and regulatory circles these days, for variety of reasons. You'll hear more about it during this conference. The AMA itself specifies stress testing that includes a consideration of how economic cycles, especially downturns, affect risk-based capital requirements. And of course scenario analysis is widely incorporated directly or indirectly in capital estimation methodologies.

Certainly operational losses can be large enough to threaten the viability of firms, and therefore could be called stressful. Indeed, tail events are in some respects the primary focus of operational risk modeling. But how does operational risk fit into stress tests that are part of an enterprise wide approach to capital adequacy assessment and capital planning? I think reasonable people can disagree quite a bit on this, and in fact I know that they do.

Enterprise level stress testing is generally scenario based. It's not a matter of thinking of every big bad thing that could happen, and calculating the impact if they all occurred at the same time. Good stress testing begins with coherent scenario development – scenarios that are severe but plausible – and carefully and thoroughly works through the ramifications of those stress scenarios, in terms of the impact on one or more variables of interest like capital or net income. So then the relevant question is, what types of operational losses would be incurred in specific stress scenarios, with those scenarios characterized by a particular configuration of driving risk factors or variables?

Some stress scenarios might explicitly include operational events, things like pandemics or hurricanes, and those are valuable tests. But many typical stress scenarios are specified in terms of macro variables like unemployment rates, interest rate spreads, property prices, and measure of economic activity. It is plausible and likely that at least some types of operational losses depend on these kinds of factors to at least some extent. But to reflect that sensibly in stress tests requires thoughtful analysis of the nature and extent of any dependence. Frankly, that's an issue that deserves more attention than it has generally received so far. It calls for data-driven analysis of the relationship between macro factors and operational losses. I also think we need to be careful about double-counting; I worry that some stress operational losses are already being captured in estimates of the impact of stress conditions on net revenues.

Conclusion

In conclusion, let's look to the future; where is the op risk world headed? I'd say look to the leaders. Leading firms recognize that there is a risk-return tradeoff around operational risk, similar to the tradeoff firms face with other types of risk. Those leading firms also recognize that in order to make decisions about that tradeoff in a sensible way, they need to rigorously measure operational risk and its sources, and therefore need state-of-the-art operational risk modeling. And finally, leading firms recognize that the use of models for that purpose – no matter how good those models are – introduces another type of risk, model risk, that must be effectively managed: the "op risk squared" of my title.

While regulators like me can encourage progress in these directions, ultimately it is people like many of you in this room who have to execute. Those of you who do so will be the success stories – and probably will be up here as keynote speakers at future versions of this conference, to tell all of us how you did it. And when you do, I'll be sitting there, paying close attention to what you have to say, just as all of you have so graciously done for me the morning.

Thank you.