Identifying, Measuring, and Managing Model Risk

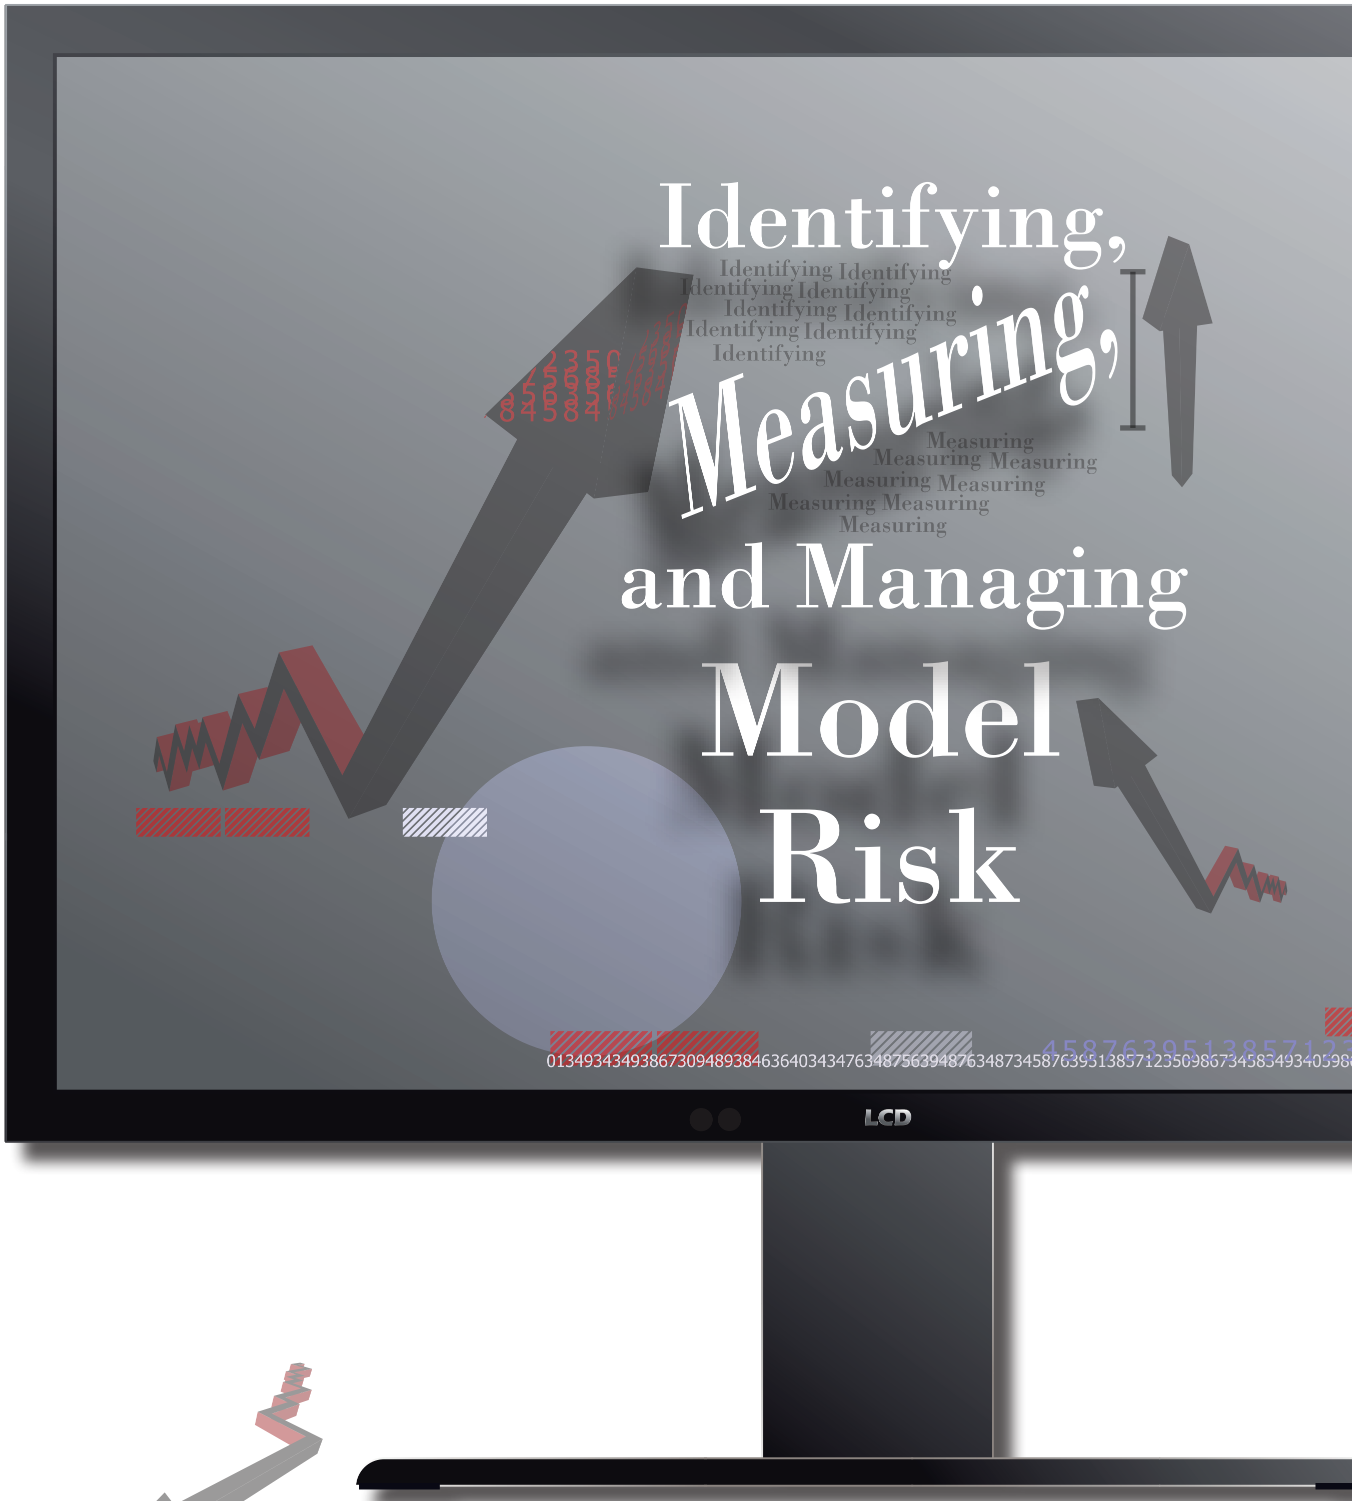

Model risk has emerged as a new field because of the challenges inherent in complex modeling. This article offers selected techniques to mitigate the attendant risks.

by Martin Goldberg and Todd Pleune

Financial companies use computer models to estimate risk, price products and investments, make decisions, and plan future strategies. However, the development and implementation of these models is a complex process that must be well controlled to ensure reliable results and protect the company's reputation and profitability. Model risk has emerged as a new field because of the challenges inherent in complex modeling.

What Is a Model?

"Remember that all models are wrong; the practical question is how wrong do they have to be to not be useful."[1]

A quantitative financial model is an approximate representation of the real world that can be used to calculate prices, risks, and strategies for financial markets. However, the real world is more complicated than any model. All financial models rely on assumptions about the behavior of the people, organizations, and other models participating in the markets. If a calculation is objectively true with no possibility of an incorrect assumption, it is not a model. Thus, the distinctive feature of a model is that it is a quantitative expression of an opinion; models are not black boxes of revealed truth, but merely numerical expressions of some view of how the world might behave. Models are a mixture of psychology, statistics, stochastic calculus, and guesswork.

What Is Model Risk?

We define model risk in finance as the risk of a loss due to a significant discrepancy between the model output and actual market experience.

This definition encompasses mark-to-model valuation, risk measurements, and any other outputs from a financial model.

Model risk is unavoidable: You can recognize it, try to measure it, manage it, and mitigate it, but almost no participant in the financial markets can get away from it completely. Measuring model risk is important, but attempting to quantify model risk has a recursive quality. Measuring model risk requires a measurement model, which itself has its own model risk.

Effective model risk management requires models that are well developed, managed, validated, and appropriately understood and used by the business. This article will analyze model risk in terms of the Protiviti Model Validation framework, with five parts:

1. Governance and Oversight—Policies and procedures governing modeling, documentation, and periodic reviews.
2. Data Inputs—Data sources and controls over data integrity.
3. Assumptions—Relevance and acceptance.
4. Analytics—Modeling theory, implementation, and testing.
5. Outputs and Usage—Usefulness and sufficiency.

Model Validation

"The saw is sharp enough if it cuts the tree."

The set of tools used to manage and mitigate model risk is called *model validation*. No model is "valid" in an objective sense; at best, a successfully validated model can be deemed "adequate" or "fit for its intended use." However, the phrase "model validation" is enshrined in many regulations, so it is used here.

Model validation is more than a regulatory reporting exercise. Proper validation of models, including the underlying assumptions, theory, implementation, infrastructure, and controls, enables the greatest competitive opportunity to:

- Derive holistic, meaningful, and accurate information for decision making, leading to improved risk-adjusted returns and loss containment.
- Enhance the business understanding of market conditions.
- Ensure performance, usage, and controls of models to meet business objectives.
- Measure consistently, and reduce reconciliation, internal model arbitrage, and ad hoc adjustment needs.
- Ensure independence and integrity of model inputs, process, and output, while limiting biases.
- Achieve compliance with various financial and regulatory standards.

Governance and Oversight—Policies and Procedures

There are several risks that start with governance and oversight. This is the overall framework in which models are developed, managed, and validated. The first step in model risk management is placing models in a framework that allows systematic model development and identification. The policies, procedures, and structures, which ensure that models are appropriate, adequately documented, and controlled, must strike a balance between being stable and well reviewed and also being living documents that match what actually happens in the field.

If policy documentation is not sufficient to determine which models to validate and who should validate them, some models may not be reviewed. This is especially likely when models are labeled "spreadsheets," "desk tools," or "calculations" and remain below the radar of corporate governance activities.

Policy should allow for discussion between the modelers, risk managers, senior management, and other interested parties when designing the models. Involvement at this early stage enhances the likelihood of eventual buy-in and acceptance of the model, and these groups may have excellent suggestions on how to solve certain model issues.

A model risk policy should outline which models require independent validation and how independence is defined. Model validation is far more effective if the validator is independent of the model developers. Having a second set of eyes looking at and approving a model gives management and the board extra peace of mind; plus, the auditors and regulators expect it. The model should not be built by the validator, and the validator's model should not be used by the developer's group.

Also, the policy should spell out whose assumptions are to be used, who has the final say in accepting or rejecting a validation, and what remediation is required for models found inadequate for their intended purpose. If a firm does not have a policy and governance framework for assigning responsibilities for model development and validation, independence will not be ensured.

Procedures must be sufficiently detailed to mitigate key-person risk. To be relied upon even in the event that its owner is indisposed, a model must be fully documented, including inputs, assumptions, analytics, and reporting. It must also include detailed operating instructions. Ideally, documentation would be sufficiently detailed to enable a knowledgeable person to rebuild the model from the documentation and get the same results.

Most regulators have a set of guidelines for model validation, which are usually quite similar. The Protiviti framework is based primarily on OCC 2000-16.[2]

Governance and Oversight—Periodic Review

In the period that Einstein was active as a professor, one of his students came to him and said: "The questions on this year's exam are the same as last year's!" "True," Einstein said, "but this year all the answers are different."

A model that has been accepted and is in use should be reviewed on a periodic basis to ensure that it is still performing as expected. Markets change, technology improves, new and better modeling techniques are devised, and the firm's views and judgments evolve, as do those of the market. Since all of these factors go into a model, revisiting the modeling and validation is a good practice and one required by most regulators.

Data Inputs

"Garbage in, garbage out."

Financial models take inputs from the observed markets and other models, perform calculations, and produce outputs. This part of the validation framework refers to the data as well as any controls of data integrity.

Data inputs are the sources of information used in the model. These inputs can be static, such as contract terms and conditions; or dynamic, such as stock prices; or in-between, such as credit ratings. Effective

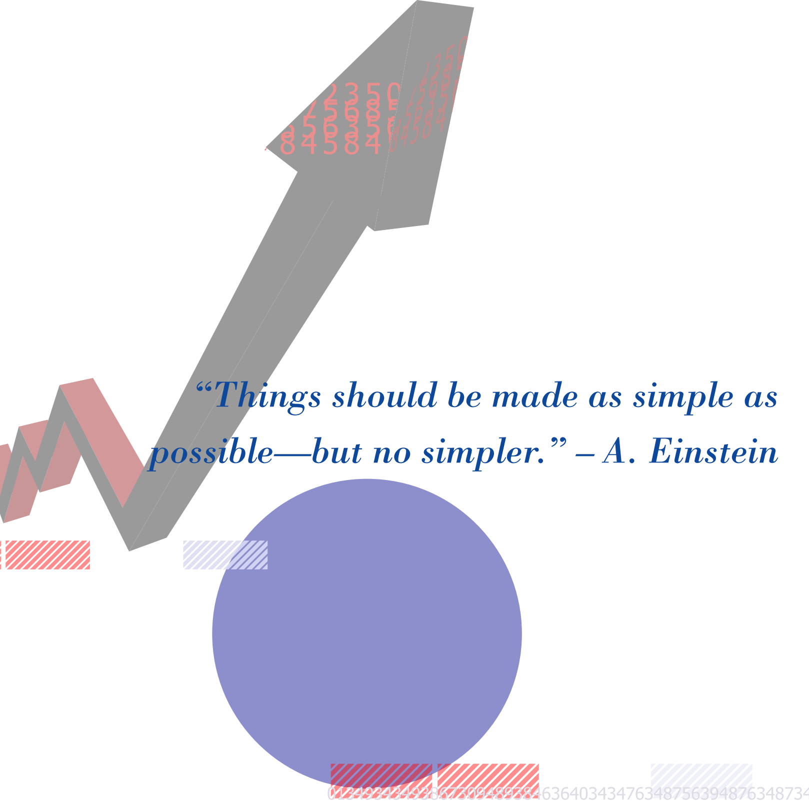

"Things should be made as simple as possible—but no simpler." – A. Einstein

modeling requires that data be correct at all times. To mitigate risk, preventative controls can be utilized—for example, developing direct data feeds so that data is not typed in by hand. Another input control is the "maker/checker" paradigm, where the checker validates the work of the maker. Detective controls, such as checksums, and other reconciliations ensure all data is incorporated in a manner consistent with the system of record. Data input risks are as simple as a security that is not properly input to the model (e.g., a call option expires on June 15, 2009, not June 15, 2008), or as complex as the input of floating point numbers into integer calculations that caused an Ariane 5 rocket launch to fail.[3] Continuous model controls and occasional model validations should check the data to mitigate these and other types of risk.

If data becomes unavailable or suspect, it may necessitate use of a different model. The risk of necessary input data becoming unavailable can be ameliorated by having a second backup data source. However, if the input is a market quote, and the market stops trading (e.g., CDO tranches of subprime mortgages), an alternative model may be required rather than just a new data feed.

Model inputs can be in error, and a "data scrubbing" process is recommended as part of any modeling effort.

Assumptions

"Investors made the mistake of assuming that housing prices would continue to rise."

More often than not, the main source of model risk lies with the assumptions. A model is, at best, only as good as its assumptions. The assumptions depend on the underlying products, the intended purpose, liquidity, convenience, and the company's preferences.

Assumptions can be classified according to who makes the assumption, whether they can be backtested, or why the assumption is being made. If the assumption is external to the firm, such as the Basel II rule that the appropriate measurement is the 99.9% one-year probability, we refer to that as a *mandatory assumption*. If the assumption expresses the firm's view on the nature of the financial market the model is intended for, we refer to that as a *market-based assumption*. If the assumption is made purely for convenience, we refer to that as an *approximation*. Finally, if the assumption is hard to explain to management, we call it a *technical assumption*.

Some examples:

- *Mandatory assumption of 80%.* "An 80% offset will be recognized when the value of two legs (i.e., long and short) always moves in the opposite direction but not broadly to the same extent."[4]
- *Market-based assumption of a floor.* "Since there is no wealth tax here, interest rates will never go below 0%."
- *Market-based assumption of a data source.* "The model will use closing prices from the Exchange."
- *Approximation of lognormality.* The Black-Scholes model assumes a lognormal distribution of stock prices because the equations are easier to work with. In most cases, the tails of the distribution are more leptokurtic (fatter), but Black-Scholes is universally the market standard quoting convention for options.
- *Approximation of a Large Homogeneous Portfolio (LHP).* The LHP approximation for CDO tranches assumes all underlying credits have the same credit ratings, financial ratios, size, etc., and that any credit downgrade will affect all of them equally.
- *Technical assumption of using a logistic regression.* In a logistic regression, the dependent variable is binary—either "yes" or "no." This is convenient if the dependent variable is a going concern/bankrupt indicator for each firm in the sample.
- *Backtestable assumption.* The firm's one-day losses will exceed our Value-at-Risk number for fewer than four days in the next year.
- *Un-backtestable assumption.* The central bank would never allow the exchange rate to fall below 150 to the dollar.

In this article, there often is a fine line between what we refer to as data inputs and what we call assumptions. Assumptions certainly include prepayment estimates, choices of proxy data for instrument valuation, and future volatilities that are estimated by the model development team or management. Hard data such as the terms of financial instruments, and other information that leads to known cash flows, are clearly data inputs. However, the choice of which data to use may be an assumption if the data does not directly describe the instrument or concept being modeled.

Assumptions must be balanced between conflicting priorities. Oversimplified assumptions limit the predictive power of the model. If jumps, skew, seasonality, or other factors are neglected, the model may not be sufficiently forward-looking in all environments. For instance, if too few stochastic processes are modeled (e.g., constant correlations for long-dated options), the model may not be accurate over time. On the other hand, too many assumptions, or those that are too stringent, may lead to a model so complicated that it is analytically intractable or takes too long to run.

An over-fitted model may trace the path of past events with no real predictive power toward future events. A significant risk with assumptions is that users of the model outputs do not understand which assumptions and simplifications were made. The assumptions should be documented in sufficiently clear language to be shared with all model users from traders to the board of directors. Only with documented assumptions will users know what to expect from a model and understand its limitations.

When a consensus exists, it is preferable to have the assumptions of a model consistent with broadly accepted market views. There is not always a consensus, or there may be a compelling reason to disagree with the consensus view. Evidence for an apparent market consensus should be documented, along with the business rationale for differing from other market participants. FASB No. 157 requires models to "reflect the reporting entity's own assumptions about the assumptions market participants would use…,"[5] so it may be necessary to have a "proprietary assumptions" model *separate* from a "market-consensus assumptions" model.

Choices about whether historical data or implied parameters are used, and whether the data is fit to a parameterized distribution or used as a histogram, are also important considerations. One of the most common assumptions is *stationarity*, the assumption that past performance is a good indicator of future results, which is the same as assuming there will not be a regime change or structural break in the nature of the market. An extreme form of this assumption, which contributed to massive failures in the dot-com sector and is ongoing in the subprime mortgage sector, is the assumption that there are no longer any business cycles and that the very recent past is indicative of the longer-term future.

Another common assumption, underlying most economic capital models, is ergodicity. This is the assumption that one year's data on 1,000 companies, or 10 years' data on 100 companies, serves as a good indicator for the next 1,000 years of one company.

Assumptions That Turned Out Spectacularly Wrong

- Latin America debt crisis, 1980s: Assumed "a sovereign nation will not default on its debt obligations." They did.
- Deutsche Bank, 1992: Assumed Black-Scholes was a good model for options—and lost about $500,000,000.
- Long-Term Capital Management: Assumed correlation was stable (Gaussian copula, no contagion); nearly lost the whole banking system. [6]
- Metalgesellschaft: Assumed abnormal conditions would revert to normal before the money ran out—abnormality lasted a year longer than their money.
- Amaranth: Assumed market liquidity would remain sufficient for orderly exit of a huge position. Went bankrupt.
- Subprime mortgage backers: Assumed the rising trend in the housing market would continue in perpetuity or at least be stable based on past performance. It didn't and it wasn't.

Analytics—Theory

"Things should be made as simple as possible—but no simpler." – A. Einstein

The theory of a model is where the assumptions are turned into algorithms and specifications. These algorithms can be intended for implementation as computer programs or manual scorecards. This involves *developmental* decisions and mathematical derivations. People using financial theory to develop and validate analytical models are called *quants*.

The theory and modeling techniques include the choices made in developing a model. Model validation should include an assessment of the appropriateness of the concepts, techniques, and rationales employed, in addition to verifying that mathematical derivations were done properly. Although the model theory can have some quite abstruse concepts and equations (this is why quants are sometimes referred to as "rocket scientists"), the validation should try to explain the theory in nontechnical language to the extent possible. The model validator must have sufficient expertise to understand and review the work of the quant developing the model.

For modeling theory and techniques, there are several areas of risk. Mis-specified market dynamics can exist if modelers do not fully understand what is being modeled. The model must be appropriate for the intended use. The number of parameters must be optimized to maximize the predictive power of the model. Too many parameters may over-fit the data or even fit noise in the data, while too few may not utilize all available information. The techniques chosen should balance the speed with which the information is needed versus the accuracy required for the model's intended use. The liquidity assumptions in data selected for the model should be reflected in the specifications.

There is interplay between developing the model's analytic theory, and the model assumptions. As developmental decisions are made, new assumptions can be made as a result. Validation should include a careful review of the rationale and evidence leading to each developmental decision. Some developmental decisions are made on the basis of tradition. For example, if a firm has a large catalog of models using Monte Carlo, it should require a strong reason to not reuse the existing Monte Carlo framework.

Analytics—Implementation

Failure to convert English measures to metric values caused the loss of the Mars Climate Orbiter, a spacecraft that smashed into the planet instead of reaching a safe orbit, a NASA investigation concluded [November 10, 1999].[7]

As with any programming effort, the implementation in computer code of the specifications for a model is a potential pitfall. This aspect of model risk also applies, but to a lesser extent, to expert judgment models such as scorecards, where the instructions and training for users in the field should be validated to ensure that the model intended to be built corresponds to the model actually being used.

Model developers should ensure that any implementation details are documented sufficiently well to be validated against the intended theory. In our experience, most implementation bugs are not due to actual coding errors, but rather to undocumented shortcuts, ad hoc decisions made by programmers using incomplete or misunderstood specifications, and other model deficiencies not found until someone codes the model.

The most complete way to validate the implementation of a model is to independently re-implement the model from the same specifications, and test the production version against the parallel version with a wide range of inputs. Any omissions or vagaries in the specifications are discovered by this process. However, because this replication is extremely resource intensive, the firm's validation policy should address what circumstances would require this form of validation, which is used less frequently in practice.

Analytics—Testing

Several different types of testing can be used to confirm model accuracy. The type and extent of testing chosen should be aligned with policy, the type of model, the historical availability of information relating to model outputs, and the criticality of decisions made with the model.

- *Backtesting* is a form of out-of-sample testing used to confirm the accuracy of models that forecast future results. A model is run using older inputs and then the outputs are compared with subsequent outcomes using goodness-of-fit tests. Once the model has been in use for a period of time, backtesting becomes an ongoing process for looking back at prior predictions.
- *Out-of-sample* testing is an alternative to backtesting, where one data set is used to calibrate a model and the model's predictive power is compared to outcomes in a separate contemporaneous data set.
- *Benchmarking* is the process of comparing the model to a "market standard" benchmark model or developing an independent parallel validator's model to test the production model.
- *Convergence testing* is the process of determining whether enough Monte Carlo runs, a fine enough grid, enough input data points for regression, or enough regression factors have been included. The model is run several times, with increasing precision, to ensure that the result is consistent and the selected level of precision sufficient. Convergence should be tested as the model is developed to ensure the model will not give erroneous results within its bounds of applicability.
- *Sensitivity testing* determines how a model's results change as the inputs or assumptions are changed. It can also provide insight into the consequences of un-backtestable assumptions on the final results. Sensitivity testing includes parameter sweeps where input values of market, assumption, or individual transaction details are varied over a wide range to get a response surface.
- *Stress testing* is used to ensure that the flexibility of the model is sufficient and that it will be stable with large changes in inputs. Failure or degradation in performance under stress indicates the limit beyond which the model is not applicable.
- *Exception testing* is the process of testing all potential financial instruments or other input data types to understand what deal terms and conditions cannot be handled by the particular model. When a deal type cannot be handled properly by the model, this exception must be documented to ensure workarounds are implemented.

The results of the tests outlined above should be confirmed using statistical tools, including:

- Goodness-of-fit tests, such as the Kolmogorov-Smirnov and Anderson-Darling tests.
- Gini curves, to compare the predictive power of different models or assumptions in a given sample.
- Monte Carlo sampling error, which should decrease with number of runs as $1/\sqrt{N}$.
- Number of Value-at-Risk exceedances compared with a binomial distribution around the stated confidence level.

Outputs and Usage

"What gets measured, gets done."

Outputs and usage represent the final part of the framework. The outputs of a model must be useful for decision making. The key questions here are the following:

- Are the outputs clear and logical?
- Are the assumptions underlying the model outputs known by the users or clearly articulated in the report?
- Are the reports regularly used to support business decision making?
- Do model users make sure they have the latest reports before making buy or sell decisions?

These questions reveal what is known as the "use test." The use test maintains that a model cannot be proved validated without being incorporated into decision making.

Business decision makers often do not have the detailed analytical backgrounds of the quants that develop financial models. This does not absolve senior management from the requirement to fully understand the assumptions, limitations, and outputs of a model used to make decisions. While this shared understanding should begin with model development and documentation, it is in reporting where clarity on the meaning and limitations of model outcomes is most important. It is essential that reports be designed to communicate results clearly and accessibly.

Regulatory Requirements

Given the potential pitfalls above, the regulators' emphasis on model validation should not be seen merely as a compliance issue, but as an opportunity to protect the company by controlling model risk. Regulators have increased their focus on the use and oversight of financial models used to support key decision-making processes. While most regulators have developed specific regulatory guidance to require appropriate governance and oversight to mitigate model risk, OCC 2000-16[8] is the first and the basis for all other such documents. OCC Bulletin 2000-16, "Risk Modeling: Model Validation," lays out the elements of sound model validation and provides guidance to help financial institutions mitigate risks arising from computer-based financial models that are improperly validated.

OCC 2000-16 described model validation requirements in terms of three high-level components: 1) an *information input* component, which delivers assumptions and data to the model; 2) a *processing* component, which contains the theoretical model and transforms inputs into estimates via the computer instructions (code); and 3) a *reporting* component, which translates the mathematical estimates into useful business information.

The guidance requires formal policies governing model validation to ensure that the level of model validation used is consistent with the materiality and complexity of the risks measured by the model. Model validators must be independent from the model developers, and they must adhere to documented policies and responsibilities.

Supervisory expectations over models can be summarized as follows: 1) decision makers understand the model's results; 2) the model is tested; 3) inputs are audited; 4) the modeling process has appropriate oversight; 5) validation is independent; 6) responsibilities are defined; and 7) change control procedures are complete.

Current Challenge—Fair Value

A major challenge faced by financial institutions is the updated requirements of FAS No. 157,[9] which apply to financial statements issued for fiscal years beginning after November 15, 2007. FAS 157 is another pronouncement in a long line of standards and regulations that place more value on accurate modeling of financial instruments. Specifically, it provides a framework for reporting the fair value of financial instruments. Fair value is the price that would be received when selling an asset or paid to transfer a liability in an orderly transaction between market participants at the measurement date. As financial instruments become more complicated, the effort needed to report assets and liabilities at fair value increases.

FAS 157 provides requirements for the methods by which fair value is reported. The "principal market" in which fair value is determined may not be the index that yields the most advantageous or positive result. A forced (duress) transaction is not a fair value. Not all inputs and modeling data are considered equally persuasive in the FAS 157 hierarchy. The hierarchy refers to the inputs to the model, not to the specific valuation techniques employed. Following are the three sources of inputs and assumptions used to model fair value in accordance with FAS 157:

- Level 1 (best) uses quoted prices in liquid markets.
- Level 2 uses verifiable data: interest rates, yield curves, volatilities, and other measures of similar instruments to infer or bracket the relevant but unobserved inputs.
- Level 3 is unobservable inputs and uses “the entity’s own assumptions about the assumptions that the market participants would use.” These can be subjective opinions, and model assumptions should be carefully stated, justified, verified, and vetted with outsiders and key stakeholders, then approved by management. Note that this is a financial-statement preparer’s guess about how “the market” would guess, which could differ from the firm’s proprietary predictions.

Determining fair values using Level 2—and especially, Level 3—data requires independent model validation. Additionally, the greater the financial reporting risk, the more a model or technique will be subject to scrutiny by auditors and examiners. The availability of inputs changes over time. The types of inputs available today may be different when the model is needed in the future. This can happen because the liquidity of a certain instrument changes over time. These factors and others require model validation to be an ongoing process where the methodology and sources of inputs are periodically revisited.

Fair value is, by nature, subjective. Additionally, FAS 157 concerns inputs to models, but is silent on some other aspects of modeling. Model validation techniques, such as sensitivity analysis and backtesting, are valuable for understanding how a model behaves with different types of inputs. Modelers should try for consistency with market consensus, and when this is difficult, internal consistency within a firm is an attainable goal.

A model suited to calculating fair value may be inappropriate for other valuation purposes if the company’s proprietary assumptions differ from those perceived as the market consensus. This disagreement is a reason one person is willing to buy and another is willing to sell.

Current Challenge—Subprime Mortgages

For many financial services companies, an even bigger issue is the effect of difficulties in subprime lending on the rates and spreads of most financial instruments. Arguably caused by a combination of models based on very recent, unproven data, as well as an assumption that housing prices would rise forever (or at least remain stable), the subprime “meltdown” has led to yield-curve shapes that have not been seen in recent memory and spreads that differ widely from historical norms. This event is characterized by the failure of more than 160 subprime mortgage companies, significant financial losses recorded by major Wall Street firms, and a dramatic decrease in liquidity.

Many companies have been forced to value certain instruments using FAS 157 Level 3 data where Level 1 data had previously been available, mainly because there are now few or no observable transactions relating to their instruments. It is a great challenge for model developers, users, validators, and external auditors when the value of marking to model rises greatly at the same time that the historical estimates of prepay speeds, PD and LGD, and credit spreads are questioned as possibly irrelevant to the 2007 market.

Given the absence of observed liquidity transactions to determine market price, there is a need for modeling as a valuation technique. It should be unacceptable to mark an instrument at par, at zero, at cost, or at the last price observed before liquidity dried up. It is not helpful to leave the position where it was priced prior to a major market shift.

Triangulation is one possible modeling framework. Triangulation models calculate a value as an estimated spread to similar, more liquid, proxy instruments that bind the value of the instrument in question. In the subprime market, the ABX index derivatives are still trading at observable prices. Similar index products are still trading in most of these newly challenged markets. If the instrument to be valued once had a relatively stable relationship to the index when both were liquid (this can be much less than the 80% correlation hurdle for hedge accounting; it’s just a best guess), then assume the previous relationship still holds, and use the previous spread to the proxy index and the current market for the proxy. This is just a starting point for modeling by triangulation, but it can be justifiable in the absence of better alternatives.

A first-principles model builds up the value of an instrument by looking at the fundamental value of underlying assets and liabilities. It requires the most quantitative skills and relies most heavily on market-based and technical assumptions. For a mortgage-backed-security model, one might use current estimates of Housing Price Index trends, prepayment speeds, loss likelihoods and severities, cash flow discounting rates, or other inputs customized to the product. A rigorous independent validation is needed to check the assumptions and how each assumption is implemented.

Conclusion

Computer models are increasingly used in banking to estimate risk exposure, analyze business strategies, and estimate fair values of financial instruments and acquisitions. As models play an increasingly important role in decision-making processes, it is critical that bank management reduce the likelihood of erroneous model output or incorrect interpretation of model results. The best defense against such "model risk" is the implementation of a sound model validation framework that includes a robust validation policy and appropriate independent review.[10]

The financial markets will continue to innovate and rely on quantitative models. All models are an idealized view of the world, and model risk is inescapable. Controlling model risk through a vigorous and thorough validation process will help prevent model risk losses and improve stakeholders' understanding of models. It is better to have a validation uncover flaws in a model than to have the regulators find the errors—or worse, have the market identify the disconnect with reality. ❖

Martin Goldberg, Ph.D., is a director at Protiviti, and Todd Pleune, Ph.D., is a senior manager at Protiviti. Contact Martin Goldberg by e-mail at martin.goldberg@protiviti.com. *Contact Todd Pleune by e-mail at* todd.pleune@protiviti.com.

References

1. Box, George E. P., and Norman R. Draper (1987), *Empirical Model-Building and Response Surfaces*, p. 74, Wiley. ISBN 0471810339.

2. http://www.ffiec.gov/ffiecinfobase/resources/retail/occ-bl2000-16_risk_model_validation.pdf.

3. http://www.cs.earlham.edu/~dusko/cs63/fpu.html.

4. Paragraph 7(v) of http://www.bis.org/publ/bcbs119.pdf.

5. http://fasb.org/pdf/fas157.pdf.

6. http://en.wikipedia.org/wiki/Long-Term_Capital_Management.

7. http://www.cnn.com/TECH/space/9911/10/orbiter.02/.

8. http://www.ffiec.gov/ffiecinfobase/resources/retail/occ-bl2000-16_risk_model_validation.pdf.

9. http://fasb.org/pdf/fas157.pdf.

10. http://www.ffiec.gov/ffiecinfobase/resources/retail/occ-bl2000-16_risk_model_validation.pdf.