Keeping all eyes on model risk

José Ramón Aragonés*, Carlos Blanco* and Kevin Dowd* look at the issues involved in model risk, considering what it entails, where it comes from and, most importantly, how to deal with it

Model risk is one of the most overlooked and underestimated risks faced by derivatives trading firms. Loosely speaking, model risk is the risk of error due to inadequacies in a pricing or risk model. Model risk is a very important subject, because inadequate attention to it can lead to very high losses. For example, if firms have poor pricing models, they are effectively giving money away to competitors; and if firms have poor risk models, they might be taking on much greater risks than they appreciate.

Some degree of error is always to be expected whenever we use a model, and we can think of this as a form of model risk. However, the term 'model risk' is more subtle than it looks: not all output errors are due to model inadequacy (for example, simulation methods generally produce errors due to sampling variation) and models that are theoretically inappropriate can sometimes produce good results (for example, simple options pricing models often work well even when they are known to be based on invalid assumptions).

Model risk is particularly large when pricing illiquid OTC contracts such as exotic positions because of factors such as unobserved variables (such as volatilities), interactions between risk factors, calibration problems, numerical approximations, and lack of transparency. Risk models are also particularly prone to model risk, because of the difficulties of aggregating risks across different positions. Such problems do not (typically) arise in 'stand-alone' pricing models.

Sources of model risk

Model risk can arise from a number of different sources.

Incorrect model specification

Model risk might arise because stochastic processes may be mis-specified. For example, pricing models that do not incorporate mean reversion or jumps will result in errors in pricing for certain markets such as electricity and natural gas. There may also be missing risk factors or mis-specified relationships. For instance, the use of constant linear correlation coefficients to price spread options can result in substantial pricing and risk errors (see box – Correlation pitfalls – below). Similarly, many energy models lead to inaccurate results because they do mis-specify volumetric variability. Model risk can also arise from mis-specified term structure or the failure to take account of transactions costs, counterparty risks, liquidity factors, or factors such as penalties in termination clauses.

Incorrect model application

Model risk can also arise from incorrect model application. A model might be poorly calibrated, or it might be correctly calibrated to begin with, and then applied in circumstances where it should not be. We might have models that can be used to price american or european options but are not applicable to other instruments such as extendible options, asian options and swing contracts. Models might also be fed poor data – "garbage in, garbage out", as the old saying goes. Data problems can arise from many sources: from the way forward curves are constructed, from data being non-synchronous, from the use of poor proxies, and so on. In illiquid markets, pricing errors can arise from the fact that the forward curves and other parameters used in the models might be based on little more than guesswork.

Incorrect implementation

Model risk can also arise from the ways in which models are implemented, and implementation requires a range of arbitrary decisions on issues such as valuation (for example, mark to market versus mark to model, or whether to use the mean bid-ask spread), the type of data to be used, whether and how to clean the P&L series, how to map instruments, and so on. But perhaps the most common implementation problems are those related to programming. In-house and third-party derivatives pricing and risk models might have errors or bugs in them (and from our experience this is the rule rather than the exception). A typical problem arises when programs are revised or modified by people who did not originally write them. Firms are also exposed to a potentially very harmful form of 'unknown' model risk where they rely on third-party pricing software and do not understand the algorithms used or have a copy of the source code.

Endogenous model risk

Another difficult form of model risk arises from the ways in which traders or asset managers respond to the models themselves (for example, how they response to value at risk (VaR) limits or VaR incentives). Traders are likely to have a reasonable idea of the errors in the parameters – particularly volatility or correlation parameters – used to estimate VaR, and such knowledge will give the traders an idea of which positions have under- and over-estimated risks. If traders face VaR limits, or face risk-adjusted remuneration with risks specified in VaR terms, then they will have an incentive to seek out such positions and trade them. To the extent they do, they will take on more risk than suggested by going VaR estimates, and our VaR estimates will be biased downwards. Indeed, VaR estimates are also likely to be biased even if traders do not have superior knowledge of underlying parameter values. The reason for this is that if a trader uses an estimated variance-covariance matrix to select trading positions, then he or she will tend to select positions with low estimated risks, and the resulting changes in position sizes mean that the initial variance-covariance matrix will tend to under-estimate the resulting portfolio risk.

Managing model risk

Given that we can never eliminate model risk, we have no choice but to

Correlation pitfalls in modelling energy spreads

Spread relationships dominate physical markets and asset hedging activities. Common examples are refinery crack spreads, power plant spark spreads, storage time spreads, and geographical spreads (shipping and pipelines, for example).

Consequently, spread options and basis or differential swaps are some of the most common instruments traded in energy markets.

Yet, despite their crucial role in most energy portfolios, the modelling of energy spreads is still very much in its infancy. Many firms still use pricing and risk models that make a number of clearly unrealistic assumptions about underlying market behaviour. In particular, many firms continue to assume that joint price distributions are lognormal or normal and that spreads are potentially unbounded, and they continue to use constant linear correlations to describe the dependence structure between related random variables.

Each of these assumptions is highly questionable. The empirical distributions for many commodity spreads clearly show that these simplifying assumptions are inconsistent with real world markets. Those unrealistic assumptions introduce significant 'model risk' that can result in serious mispricing of spread derivatives or the mis-measurement of the risks they entail.

Market practitioners often respond to these issues by using different correlations for different purposes: they use correlations with ad hoc adjustments for time, maturity, different shocks or different strike prices, and so forth. However, the fact remains that the two-factor model assumes that the 'true' correlation is constant – and this highlights the deeper point that correlation is not an adequate measure of dependence in these sorts of problems.

Far from being the all-purpose dependence measure it has often been assumed to be, the (linear) correlation coefficient can only be used as a measure of dependence between risky variables that are elliptically distributed or are subject to elliptical shocks, and this applies to few (if any) energy or commodity price processes.

If we can't use correlations to measure dependence, then what should we use? Perhaps the best answer is to use copulas. Copulas are an alternative dependency measure that is reliable when correlation is not. Copulas enable us to extract the dependence structure from the joint distribution function, and so 'separate out' the dependence structure from the marginal distribution functions. This is extremely helpful. For energy and commodity price modelling copulas allow us to model individual asset price behaviour using non-normal processes (such as a mean reverting jump-diffusion), and still have a 'correct' way to determine their joint behaviour. We can do so by taking marginal distributions – each of which describes the way in which a random variable moves 'on its own' – and then applying a copula function that tells us how the marginals 'come together' to determine the joint or multivariate distribution. ❑

Who should check the validity of the risk models?

A partial answer, inevitably, is that much of the responsibility for vetting risk management models must lie with the risk management function itself. In larger firms, one could envisage this task as carried out in a separate unit within the risk management function, so one unit of risk management builds risk models, and the other checks them. In addition, it is also good practice for firms to institute periodic risk audits – audits of all aspects of the firm's risk management, including audits of the models used, carried out by external parties with the skills to do so[1]. Such audits would also provide useful feedback, especially on areas where the firm was lagging behind best market practice. ❑

[1] Such audits are not only good practice in themselves, but also help senior managers establish that they have carried out due diligence – which is a particularly important consideration for firms operating in the US with the advent of Sarbanes-Oxley and other corporate governance reforms.

learn to live with (ie, manage) it. There are many ways it can be managed, and these fall under three main headings: those applicable by the individual risk practitioners who build models and use them; those applicable by the managers whom the risk practitioners report to; and institutional methods that involve the establishment of suitable institutional or procedural structures to detect and counteract model risk.

The people involved in the risk function need to be intelligent, have integrity and have good experience of the institutions and markets in which they are operating. In addition, the risk team must also include people who have a strong knowledge of statistics, derivatives pricing, modern portfolio theory, and risk measurement, experience in energy and commodity derivatives trading and/or risk management as well as good knowledge of spreadsheets, relational databases and programming languages (for example, VB, C/C++/C#, MATLAB, Mathematica).

Senior managers should have some basic appreciation of the strengths and weakness of the models used, and of how the models are being implemented. They should be able to understand what their risk managers are talking about, and should have a good appreciation of how pricing and risk estimates are affected by trading strategies (ie, they should be aware of potential endogenous model risk). More fundamentally, senior managers need to be on their guard against the tricks that people play – how traders can hide losses and 'game' VaR models to their advantage, and so on. In combating these sorts of problems, managers also need to recognise that if they are to be effective, their risk managers must have the knowledge and skills to match traders, and this means that they must be suitably remunerated – otherwise, there is little incentive for those with the necessary skills to want to do anything but trade.

Senior managers can also reduce their vulnerability to model risk problems by encouraging a multi-disciplinary team approach to model building. They should not see models as incomprehensible formulas that quants or risk modellers hand over to programmers to make even more incomprehensible. Instead, they should see models as the product of an inter-disciplinary team, involving inputs from mathematicians, statisticians, computer scientists, finance experts, accountants, traders, model users, and others. They should encourage people from these disparate groups to understand each other, and foster a climate of constructive criticism.

With regard to institutional methods, first and foremost, the right governance structure needs to be established to protect the independence of the risk management function. Within this governance structure, one basic defence is a sound system to vet models before they are approved for use and then check and periodically review them. The goal of model validation is to avoid situations where models are inadequate for the problem at hand, or are applied improperly. Model validation should be conducted in conjunction with back testing and price verification in order to understand the sources of model risk. All this work should be carried out free of undue pressures from the front office, and traders should certainly not be allowed to vet their own pricing models.

Even though best practice risk management principles call for firms to keep a database with all models used to price derivative instruments, as well as the methods used to calibrate model parameters, it is rare to see that information in a centralised location that can be accessed by senior management and audit groups.

Documentation for specific models should be sufficiently adequate to facilitate independent reviews, training of new staff, and detailed enough to allow exact replication of the model should key staff move to other areas or depart the company. At a minimum, this database should include: summary overview of general procedures used; reasons for choosing those procedures; descriptions of model applications and limitations; records of key personnel involved and milestone dates in model construction; and descriptions of validation procedures and results.

A key plank of the firm's governance structure is independent risk oversight (IRO): the firm should set up an IRO group or middle office unit. This unit should encompass risk measurement as well as risk management,

Charging for model risk

Besides having good model vetting procedures and independent risk oversight, there are also other ways that firms can manage model risk. One sound practice is for them to keep reserves against possible losses from model risk. The reserves attributed to a position should reflect some measure of the model risk involved, so that positions with higher model risk get higher reserve-charges than positions with lower model risk. Such charges not only provide the firm with a cushion to absorb possible losses from model risk, but also help to ensure that the cost of model risk is accounted for and attributed to the positions concerned. Firms can also counteract model risk by taking account of it in setting position limits: thus, if a position is known to have considerable model risk, a firm can limit its exposure to this source of model risk by imposing a tighter position limit. ❑

❑ model risk

should be independent of line-execution areas (such as treasury, trading, portfolio management and asset-liability management), and its head, the chief risk officer (CRO) should report to the CEO and, ideally, sit on the board or other governing body. The middle office should have a clear mandate from senior management, and its policies should reflect the corporate policies towards risk – the corporate risk appetite, and so on. It should also have its own designated budget and, to help mitigate any temptation for the middle office to go along with excessive risk-taking elsewhere in the organisation, the remuneration of the CRO and his or her staff should not be tied to the performance of other units (such as trading profits).

This unit should have the authority to block any trading or asset management activity within the organisation, and have the authority over the use of all pricing or risk models. It should have responsibility for monitoring risk independently of other business units, and particularly the front office. In carrying out these tasks, it should seek to ensure a balance between an excessively prohibitionist stance on risk (ie, everything not expressly allowed is forbidden) and an excessively lenient stance (ie, everything not expressly forbidden is allowed), and also aim to ensure that all interested parties are fully involved in the firm's internal risk measurement/management dialogue.

The middle office should also be responsible for all aspects of risk estimation, including stress testing, back testing and (at least some) contingency planning, for ensuring that all models are adequate for the tasks to which they are being used (ie, take responsibility for vetting, checking and monitoring the models used), for reporting and disseminating risk information throughout the organisation, and for protecting and maintaining the integrity of the firm's risk measurement and risk management systems.

A management report covering model risk such as the one in Figure 2 can highlight the areas where model risk can have some material impact as well as the action plan to investigate and resolve the main issues.

To summarise, model risk is one of the most important and least appreciated areas of market risk management. Firms are particularly vulnerable to its dangers because of its insidious and hidden nature, and because it is all too easy for us to start 'believing' our models and in do doing lose sight of their limitations. These mental habits then make us vulnerable in those situations where reality rudely asserts itself and reminds us that our beliefs can sometimes be disastrously off the mark.

Model risk therefore casts its shadow over everything we do in risk management. The only real defence against it is to keep on our guard and keep asking ourselves what would happen if our assumptions fail to hold. Of course, this is easier to say than to do, but the financial markets are littered with the remains of those who have ignored model risk and thought they could get away with it, and these include some of the biggest names of their day. At the end of the day, model risk is like the proverbial ghost at the banquet – an unwelcome guest, but one that we would be very unwise to ignore. ❑

Figure 1.
Ten ways practitioners can protect themselves against model risk

- ❑ Be aware of model risk
- ❑ Keep a central database to document models used for pricing and risk purposes
- ❑ Identify, evaluate and check key assumptions
- ❑ Test models against known problems
- ❑ Choose the simplest reasonable model
- ❑ Back test and stress test the model
- ❑ Estimate model risk quantitatively
- ❑ Don't ignore small problems
- ❑ Plot results and check summary statistics (ie, mean, standard deviation, min, max)
- ❑ Re-evaluate models periodically.

Source: Dowd (2005)

Figure 2.
Management Model Risk Report – August 2005

Dashboard Indicator

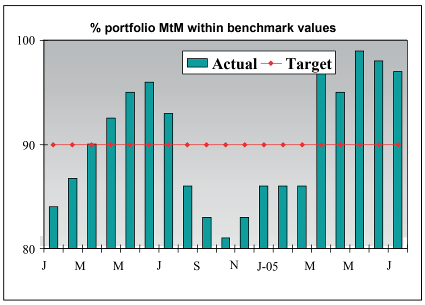

Current MtM Discrepancies by Desk		
Desk	**MtM**	**Difference**
Natural Gas Trading	657,890	-4%
Oil Trading	-435,665	5%
Power Trading	2,765,980	12%
Soft Commodities	-187,654	7%

Current MtM Discrepancies by Counterparty		
Counterparty	**MtM**	**Difference**
JP Morgan	-345,875	2%
BP	547,896	1%
Shell	345,980	19%
Duke	-675,450	3%
Sempra	1,234,456	12%

Main Problems detected
Valuation of OTM Natural Gas options. Volatility smile discrepancies.
Portfolio of long term OTC electricity forwards overvalued vs. counterparty values.
Large MtM discrepancies in positions with Shell and Sempra

ACTION PLAN		
Activity	**Who**	**When**
Research alternative option valuation models and present to risk committee	Anne	8/15/2005
Validate volatility surfaces for vluation of OTM options	Kim	8/15/2005
Reseach source of discrepancies in value of positions with Sempra	Kim	8/15/2005

Additional Comments: Since March, discrepancies in MtM are above threshhold levels. Risk group is investigating sources, but the source is yet to be found. Out of the money options are the instruments with larger discrepancies. Research into alternative models and volatility surfaces will likely shed some light into the sources of the problems.

Source: Black Swan Risk Advisors

***José Ramón Aragonés** is professor of finance at the Universidad Complutense in Spain (aragones.jr@ccee.ucm.es), ***Carlos Blanco** is MD of Black Swan Risk Advisors (carlos@blackswanrisk.com) and ***Kevin Dowd** (kevin.dowd@nottingham.ac.uk) is director of research at Black Swan Risk Advisors and professor of financial risk management at Nottingham University Business School in England