

Risk

JULY 2001
www.risk.net

CURRENCIES • INTEREST RATES • EQUITIES • COMMODITIES • CREDIT



Markets and models

Emanuel Derman, Managing Director, Firmwide Risk, Goldman Sachs, & Co.

Address: 10 Hanover Square, New York, NY 10005.

Telephone: 212-902-0129

Web: emanuel.berman@gs.com

risk waters group



Markets and models

Emanuel Derman, managing director, firm-wide risk at Goldman Sachs, originally wrote this primer on how risk managers help price derivatives for his firm's traders. We felt it deserved a wider audience

Trading desks in many product areas at investment banks often have substantial positions in long-term or exotic over-the-counter derivative securities, custom-tailored to satisfy the risk preferences of their customers. The idiosyncratic nature of these securities makes them relatively illiquid.

As a consequence, these positions are marked and hedged using sophisticated and complex financial models, implemented as software and then embedded in front-office risk management systems. These models derive prices from market parameters (volatilities, correlations, prepayment rates or default probabilities, for example) that are forward-looking and should ideally be implied from the market prices of traded securities.

Because of their illiquidity, many of these positions will be held for years. Yet their daily values affect the short-term profit and loss of the banks that trade them, and also inevitably influence the careers and compensations of the traders and salespeople who structure, trade, hedge and market them.

Increasingly, therefore, banks employ autonomous firmwide risk managers,

situated several steps away from the trading desk, who represent the firm as a whole rather than a particular desk, and whose mission it is to verify the marked prices of these positions.

Derivatives, models, prices

For liquid securities, price verification is straightforward; one simply checks that the prices assigned by the trading desks correspond with those obtained from reliable external sources. Illiquid derivative securities, our focus here, are more complex.

A derivative security's payout is determined by how it depends on the prices of its underlying securities. So, finding the current value of a derivative security is a question of relativity: what is the derivative security worth relative to the known prices of its underliers, given their estimated future behaviour?

The dependence of the payout on the underliers' values is determined solely by the contract, and should be unambiguous. But the estimated future stochastic behaviour of the underliers is a matter of opinion, which is where models enter.

Most derivative securities are non-linear: their prices vary disproportionately to

the prices of their underliers. Traders need models to augment their intuition, in order to interpolate from the linearly changing underlier prices to the non-linearly varying prices of the derivative securities. A model tells you how to manufacture the derivative from its linear underliers. Options theorists and practitioners like to say "replicate" or "hedge" instead of manufacture, but the meaning is identical.

You can think of a derivative as a mixture of its constituent underliers, much as a cake is a mixture of eggs, flour and milk in carefully specified proportions. The derivative's model provides a recipe for the mixture, one whose ingredients' quantities vary with time.

The derivative's model or recipe tells you how to dynamically vary the proportions of the mixture as time passes and the underlying security prices change. If you have estimated the future stochastic behaviour of the underlying securities accurately, then, as their prices change and time passes, the mixture's price will behave identically to that of the derivative security itself.

Derivative models work best when they use as their constituents underlying secu-

rities that are one level simpler and one level more liquid than the derivative itself. Thus, one manufactures a standard stock option out of a changing mixture of stocks and riskless bonds. Similarly, one manufactures an exotic option, for example a knockout call, out of standard options.

In any model, the value of the security and its hedge ratios (the proportions of constituents in the mixture) are intrinsically related to each other – obverse sides of the same coin.

Any model for the future behaviour of an underlier is an attempted simplification of a reality that is likely to be much more complex and unpredictable than the model itself. Therefore, it is not obvious which model is correct, nor even exactly what correct means. Also, once you've picked a model, the input parameters you use to describe its details are themselves estimates of future security characteristics, such as volatility or correlation, whose exact values are unknown.

When you create and use a model to value a derivative security, you need to:

1. obtain a careful description of the security's payout in terms of the values of each of its underliers;
2. specify a model for the future behaviour of the underliers;
3. obtain accurate values of model parameters (volatility, for example) that describe the behaviour of the underliers;
4. calibrate the model to these parameters so that it reproduces the known prices and assumed future behaviour of the underliers; and then (usually)

5. build a computer program that incorporates items one through four above.

Derivatives' price verification

Simple derivatives, such as standard stock or index options, usually have well-accepted models that are used as a market-quoting convention for communicating price information.

Market participants who use a broadly accepted model are quoting their prices in terms of parameters (such as volatility) that, once entered into the model, produce the quoted price as output. Their use of the model does not necessarily imply that its assumptions about the future behaviour of the underliers accurately reflects their actual behaviour, though some correspondence between assumptions and behaviour must obviously exist for the model to be embraced.

For liquid derivative securities, one verifies a marked price by checking that the trader's inputs (such as volatility) agrees with those obtained from an observed external market. For less liquid or more exotic derivatives, one must verify the payout description, check the plausibility of the inputs and examine the reasonableness of the model itself. One must also test the computer programs that incorporate the model.

To do all this, one should ask the following series of open-ended questions:

1. Does the model embody an accurate description of the terms of the derivative's payout? The function that specifies the payout can be exceedingly complex,

involving, for example, subtleties of price averaging, barrier crossing, holidays and date conventions. In practice, many of the errors made in marking derivatives involve accidental or sloppy mis-specification of these details in the trading system's database.

2. Does the model provide a realistic (or at least plausible) description of the evolution of the factors that affect the derivative's value? When a trader books a derivative transaction into a firm's risk system, he implicitly assigns a valuation model to it. All models are simplistic. Is the assigned model an appropriate simplification? In practice, many of the errors made in marking derivatives are the consequence of an unsuitable assignment, or perhaps an assignment suitable only during the current market regime. For example, an adequate model for a hybrid product that incorporates both equity and currency risk may often justifiably neglect exchange rate volatility during normal, relatively low-volatility times.

3. Has the model been accurately calibrated to the observed behaviour, parameters and prices of the simpler, liquid constituents that comprise the derivative?

4. Finally, is the software reliable? Handling portfolios of complex derivatives requires building intricate and extensive programs and systems that will inevitably contain errors. Accurate valuation demands an organised

Risk management procedures explained

Because the use of models requires an intricate interaction between mathematics, software, systems and common sense, it is useful to have procedures in place for the validation, modification and maintenance of models produced in each product area.

Model control

All new derivatives models should be tested and certified by someone other than the developer. Alterations to models and software should be logged and documented.

Quantifying a desk's model risk

Each product area should use its front-office risk management system to report the absolute value of the sensitivity of each position in their portfolio to volatility, interest rates, or other key determinants of derivative value. These sensitivities can be summed and grouped into bins, with each bin corresponding to one of the range of models used to value derivative products. Ranking the bins by decreasing sensitivity directs attention to those models whose impact on the desk's risk is greatest.

Model validation

The models responsible for the greatest sensitivity of the portfolio should be documented and validated. The developers and users of a particular model should produce a written report demon-

strating that the model captures the relevant features of both the market and product it addresses, and should show that it has been adequately designed, built and tested. The report should also address the numerical accuracy and stability of the algorithms and software, and emphasise those domains where inaccuracies or approximations become unacceptably large. This documentation will provide a long-term corporate memory of the principles and implementation of the models.

Rational documented model assignment

When a trader books a deal and assigns a valuation model to a complex OTC derivative security, there should be an explanation of why the particular model was chosen, the simplifications it involves, and why they are currently justifiable.

Input testing

The simplest test of prices is to report the extent of the correspondence between the input parameters in the desk's models and the same parameters obtained from external market sources. The fraction of the desk's sensitivity to its key derivative parameter (say, volatility) that has been successfully compared with external sources is a useful, if somewhat naive, quantitative measure of the scope of verification.

However, input testing can be complex. Some parameters are

not single numbers, but one-dimensional curves of interest rates, or even multi-dimensional surfaces of implied volatilities.

Some more complex input parameters (for example, default correlation, spread volatilities or default recovery rates) have no reliable external market sources because of the lack of liquid trading instruments. For these, one must estimate their most reasonable implied values indirectly, by means of theoretical models that link their values to those of other traded instruments.

Comprehensive price verification

Nothing is better than a completely independent check of price and hedge ratios. A strategist knowledgeable about the market, but organisationally separate from the trading desk, should start with the confirmed trade details and build an independent model to describe product and the market, then calibrate the model, and finally provide an estimate of value and hedge ratio. This is the most thorough and valuable way to check prices. It provides a truly autonomous verification of value and, equally important, highlights the sensitivity of price and hedge ratio to the choice of model.

Valuation adjustments for transaction costs, hedging error and model risk

It is never clear what profit or loss will result from hedging a derivative security to its expiration. Markets will move in unexpected ways, sometimes intensifying transactions costs and often dismantling what may have seemed a reasonable hedging strategy.

These effects are rarely captured by the conventional models used in front-office trading systems. Therefore, for illiquid positions, it is important to estimate the adjustments to the conventionally marked values that can occur as a result of long-term hedging.

Monte Carlo models can simulate both underlier behaviour and a trader's hedging strategy, to create distributions of the resultant profit or loss of the entire portfolio. These distributions can be used to determine a realistic adjustment to the trading desk's conventional marks that can be withheld until the trade is unwound and their realised profit or loss determined.

These adjustments for hedging vary, as they should, as an option moves in or out of the money, or as volatility changes. Monte Carlo analysis offers a good sense of the variation in portfolio value that will

be exhibited over the life of the trade due to transactions costs, hedging error and model risk. Ultimately, such analyses should be part of the desk's own front-office valuation system.

Periodic comprehensive model review

Immature derivatives markets often display prices that are consistent with the usual Black-Scholes assumptions. Then, as markets mature and participants gain experience of the supply, demand and shocks that their underliers and derivatives can experience, prices start to reflect these realities more accurately.

For example, ever since the 1987 stock market crash, equity index derivatives have displayed a skew in which out-of-the-money puts trade at much higher Black-Scholes implied volatilities than out-of-the-money calls. During the past 10 years, skews of this type have become prevalent in most derivatives' markets, from swaptions to gold options.

Skew is inconsistent with most simple and widely used options models. A major challenge in determining the accuracy of illiquid exotic derivatives prices is the battle to develop realistic models that can be calibrated to reflect the skew of standard options, and which can then be used to value the exotic ones. For listed equity index options, liquidity is good and the standard skew is easy to observe. In swaptions markets, accurate information on the skew is much more difficult to obtain.

It is therefore generally advisable to periodically revisit entire derivatives markets and their models, to examine existing approaches and perhaps develop new ones.

Since it is never clear exactly which model (for example, for the skew) is correct, it is a good idea to investigate the effect on pricing and hedging of a variety of plausible models that can be calibrated to the market, and so understand how sensitive the desk's marks are to the particular model they use, and the assumptions it makes about the future.

Pre-review is better than verification

Price verification suggests a check done after the deal is closed, but the most effective way to verify prices is in advance, in order to provide independent analysis and confirmation on pricing and hedging before deals close. Divisional risk managers and desk heads should seek to have the prices, hedges and risks of complex deals reviewed by independent analysts in advance of closing. Assurance in advance is better than verification (or lack thereof) afterwards. ■

framework for the control and testing of the development, alteration and release of models and code.

A multi-faceted approach to price verification

A derivative security that cannot be bought or sold in a liquid market has no unambiguously correct value. Since it may have to be held and hedged until it expires, its realised value will depend on the future behaviour of its underliers and their markets, as well as on the details of the hedging strategies adopted by the traders managing the portfolio.

A firm's first, broadest and most efficient line of defence is the quality of the traders, model builders and software developers in each product area. Ultimately, one relies on them to be careful, honest and responsible.

In particular, education is important;

traders cannot be too knowledgeable about the assumptions behind models and their limitations. Nevertheless, it is wise to provide independent price verification, delivered from the vantage point of a firm as a whole, whose interests may not always coincide with those of an individual trader.

Some of the methods recommended in the box on page 50 are derived from practices developed within the Derivatives Analysis group in the firmwide risk department at Goldman Sachs. This group collaborates with controllers and the divisional and firmwide risk managers to provide independent price and model verification for selected complex or illiquid derivative securities.

The group consists mostly of quantitative derivatives strategists, similar to those employed in building and using models for derivatives trading desks. Each member

of the group specialises in a particular product area, ranging from swaps through equity, currency and commodity derivatives, to more recently developed markets such as credit derivatives. Importantly, the group also contains several ex-traders with a strong quantitative background who bring familiarity with markets and their practices.

Because of the subtleties of models and the unattainability of "correctness", a multi-faceted approach to price verification is best. One should selectively span the valuation process at all levels of detail. The methods described in the box cover a broad spectrum, moving progressively from coarser to finer scales of resolution. ■

Emanuel Derman is grateful to Bob Litzenberger, Noel Donohoe, Slim Bentami and Leon Tatevossian for insightful advice and remarks