

REPRINTED FROM

Risk

RISK MANAGEMENT | DERIVATIVES | REGULATION

www.riskcare.com

MAY 2010



Models

Towards a global valuation model

by Claudio Albanese, Guillaume Gimonet
and Steve White

 RiskCare™

Towards a global valuation model

Banks use a variety of pricing models across business lines, creating discrepancies in the way various financial instruments are priced. But developments in high-throughput computing could lead to the possibility of a global valuation model, argue Claudio Albanese, Guillaume Gimonet and Steve White

Every bank uses pricing models that are specific to the instrument it is valuing. But even when they are perfectly implemented, there are inconsistencies between the models used for each instrument on a firm-wide basis.

Every valuation formula embodies two common weaknesses:

- Each explicitly or implicitly makes assumptions about the future states of underlying risk factors. Therefore, the modelling of the future is specific to each individual instrument, rather than to each risk factor, and is not standardised across the whole firm.
- Each model is calibrated to a set of chosen observable market prices. However, the values used for calibration are limited and instrument-specific, rather than being broadly representative of the market.

While each position is ‘correctly’ calibrated and priced in isolation, this instrument-centric approach to modelling creates inconsistencies in the valuation process that affects hedging, the balance sheet and capital adequacy provisions from a broader, firm-wide perspective.

This is not model risk in the traditional sense, which is typically taken to mean incorrect specification, usage or application. Instead, the underlying assumptions on which the mathematical representations of the models are based cannot be reconciled from instrument to instrument and from model specification to model specification.

These inconsistencies are widely accepted, but are they reasonable and inevitable, and does inconsistency necessarily mean there is error? The quantitative finance equivalent to the grand unified theory in physics is global valuation, where a single model can be used to price all instruments – but does such a model exist?

According to the fundamental theorem of finance, prices obtained by consistently using a single model are guaranteed to be arbitrage-free. Conversely, given a set of arbitrage-free prices, there must be a single model that can reproduce them (otherwise, by definition, the prices will allow for arbitrage). So the model must exist – but can we find and implement it?

Such a model has not been computationally possible until now, but with the arrival of high-throughput computing (HTC) – the generalised application of graphical processor units (GPUs) – the amount of processing power available for mainstream industrial applications at commercially attractive rates has dramatically increased.

However, GPUs do not work like central processing units, so the method of pricing financial instruments needs to be redesigned if it is to harness this computational power.

The core of valuation theory centres on the need to compute transition probabilities for a given model specification. Theoretically, these transition probabilities are the result of repeated matrix multiplications, but a brute-force approach has been computa-

tionally out of reach. Many classical mathematics tools, such as special function theory, group theory and spectral theory, have been developed to work around this problem, and these techniques have led to mathematical finance as we know it today.

The reliance on analytical tractability imposes a straitjacket on modelling possibilities, limiting the flexibility and comprehensiveness of models and their ability to achieve economically realistic representations. The original justification for these techniques was that they could be realised using the technology of the day. With HTC, this logic needs revisiting.

By shaping the problem to the tool, the power of the technology can be efficiently exploited, and this favours numerical approaches (simulations) over closed-form mathematical models.

An HTC-friendly solution would solve both problems above – consistent modelling of future states and broader calibration. This can be achieved by restructuring the valuation process flows.

To do this, existing pricing models need to be decomposed into three discrete process flows:


- the simulation of potential future states of risk factors;
- valuing each instrument in each future state and aggregating into groups (by book, counterparty, etc); and
- post-processing these results to determine fair value, exposures and sensitivities.

By restructuring the valuation process in this way – and in particular, by making the first and last of these components generic across all instruments – it is possible to create a valuation framework that treats all financial instruments consistently. There is no logical or valid reason why the risk factor calibration needs to be repeated for each instrument – the fact this practice is nearly universal is a consequence of the evolution of the financial industry.

Not only does this decomposed approach enable centralised and therefore globalised risk factor modelling, but it also provides the springboard for broader calibration: rather than ‘fit’ individual products to the market, the risk factors themselves in their simulated states are fitted to a broad spectrum of observable market activity.

Figure 1 depicts the generic change in process flows between the two approaches. At the highest level, all

1 Changes in process flows between instrument-specific and global valuation techniques



In instrument-specific valuations, the risk factors are mapped into the instrument prior to the simulation of future states. For global valuation, the simulation has to be performed to the risk factors prior to the instrument mapping.

Key:

Processes are in italics, data structures are square boxes.

RF = risk factors.

INST = individual financial instruments, portfolio positions, aggregated positions (eg, by counterparty)

valuation systems are the same – a set of inputs are fed into a pricing function that produces a set of outputs. The inputs take the form of current and historic market data and counterparty information, and the outputs are values, sensitivities and exposures.

The bulk of the functionality in the pricing algorithm is removed from the instrument-specific valuation, so only the

to meet modelling objectives and to define the problem in such a way that it can leverage the computing capabilities of HTC.

The description of the methodology in the box overleaf presents a high-level approach in order to give an overall perspective on the solution. It is not intended to trivialise the complexity of the undertaking.

The reliance on analytical tractability imposes a straitjacket on modelling possibilities, limiting the flexibility and comprehensiveness of models and their ability to achieve economically realistic representations

payout logic remains. Instead, most of the analytical and computational effort goes into simulating each risk factor. A ‘risk landscape’ is created, where every possible future state of every risk factor is computed and is probability weighted.

Implementing global valuation

There is no single right way of implementing a global valuation model – the framework described here is intended

Instrument-specific valuation – the way the industry works at present – has come about with good reason, and may be considered a victim of its own success. The context in which models originated was instrument-specific, and it is only because of the current scale of the practice – firm-wide rather than desk-specific – that consistency has become a problem.

From the technology perspective, closed-form solutions provide a low entry

point for developers. They can be implemented with relatively simple (re-entrant or non-thread-safe) programming, and there is a low dependency upon infrastructure (all you need is a spreadsheet). From a management perspective, the separation between quants and IT is well understood – projects are easy to structure into work streams and milestones, and do not require large-scale resourcing. But we need to be prepared to lose this if we adopt the approach described above.

Since HTC computing in effect means starting over again, there need to be some very good reasons for doing so. Not only are the models that ‘work’ thrown out, but also a lot more labour is also being taken on. So, what are the business benefits that can be expected from global valuation?

■ **For senior management.**

■ *A second view.* An economically consistent model offers an alternative viewpoint for valuation and arbitrage checking – valuations obtained from a global valuation system can be compared with those obtained using existing systems. It is unlikely these discrepancies will identify profitable trading strategies, as the differences exist because of accepted market practices. More likely, this information will identify pricing imbalances and bubbles. This is practical and, in the right hands, provides very valuable information – irreconcilable data ought to give advance warning to senior management of looming problems within their business.

■ *Capital deployment.* Capital allocation can be economically targeted. By centralising risk factor calculations, investment and exposure limits can be more readily expressed in terms of risk factors. Senior management can then take a view on how to focus the balance sheet, hedge out undesirable risks, and apportion capital to accentuate those risk factors/products/counterparties that are able to achieve the highest risk-weighted returns.

■ **In the front office.**

■ *Risk-informed decision-making.* Distributing the risk landscape cubes to traders enables them to make cost-benefit decisions in terms of economic capital, rather than just expected profit and loss. If risk quantification is standardised and distributable, then it can facilitate economic capital on the desktop. All trading decisions will treat capital in an equal manner.

■ *Freer modelling.* Economists and traders can create models embedding their own views and techniques by designing process dynamics. These can then be executed in the global valuation engine with no detriment to computational performance.

■ *New functional possibilities.* Performance gains for firm-wide revaluations can open up a range of applications that have been impenetrable so far – pre-deal credit checking, value-at-risk-based margining, intra-day full revaluation and real-time risk decomposition. Also, since global calibration provides a forward-looking analysis rather than a historical perspective, it may open up uses such as asset allocation, credit rating assessment and forecasting.

■ *Faster innovation.* Centralised risk factor modelling gives a safety-first approach to financial innovation – the risk modelling is more comprehensive and better controlled, and the barriers to new product introduction are significantly reduced.

■ **In risk management.**

■ *Applied risk management.* Global valuation will be an enabling technology for better risk management – macro hedging will be more accurate, and therefore cheaper and more efficient. This is the level at which the broader calibration effects come into play.

■ *Integrating front-office valuation, market and credit risk.* The global valuation methodology has striking similarities with the simulation frameworks that many market and credit risk systems deploy in terms of process flows. All the measures credit and market risk systems rely on are by-products of the global valuation framework. It is still necessary to integrate this computational framework into credit and market risk workflows and functionality, but it is important to be aware of these eventual end-uses from the outset, so the technology is designed with appropriate provisions.

■ **Better economies of scale.**

■ *Organisational effects.* Taking a holistic, firm-wide approach to valuation, as opposed to an instrument-by-instrument approach, actually reduces the sophistication and volume of the maths. Instead, the demands made upon technology and management have increased, but this is spread across products and functions throughout the organisation.

■ *More efficient coding.* Polymorphism can be exploited by maintaining a single computational engine with core functionality shared across product lines and asset classes. By relaxing the constraint of analytic solvability, the technical differences between rates, foreign exchange, equity and commodity modelling are reduced, commonalities can be shared in base classes, and the differences can be implemented in derived classes.

An impact analysis

Since those banks that are currently working on global valuation are doing so on a limited basis and with varying objectives, there are no established methodologies for realising it. However, some of the likely organisational effects can be anticipated.

■ *Power of executive management.*

Executives will get more control since risk is more visible. It will also allow them to overlay the control variables of product and sector with the risk factors a bank wants exposure to.

■ *The rise of risk.* The most significant aspect of organisational change presented by global valuation is that risk factor modelling necessarily becomes a centralised activity, and therefore ought to be owned by the chief risk officer. The chief risk officer moves from being a controller to a driver, and a creator of value rather than a business overhead. This profoundly changes the motivation for investing in risk competency and technology. Rather than being seen as a cost, risk management becomes a contributor to profitability, requiring investment.

■ *The rise of technology.* The emphasis in expertise changes from the dominance of applied mathematical methods to a more evenly balanced dependency upon technological excellence – the techniques of computer games developers become critical to investment banks. Building bespoke systems will require world-class industrial-scale management competency, similar to that in the aviation, energy and pharmaceutical sectors.

■ *New products group.* Introducing new financial products becomes much less effort – it is simply a case of ensuring the payout description is accurate, then applying this to pre-existing risk factors. A full revaluation model no longer needs to be developed, distributed and validated – modelling and testing the payout is less effort. The bottleneck for

Methodology

Step 1: Risk factors initial set-up

For each risk factor, a matrix of discrete states is created. For example, an overnight interest rate may have values between 0% and 20% in increments of 25 basis points.

Time is divided into epochs within which time-dependent parameters are held constant.

A transition probability matrix is generated based on the initial risk factor process parameter estimation for each epoch. Transition matrices for each epoch are generated by multiplying the single-step transition matrix.

The result is a collection of three-dimensional matrices – one per risk factor – but at this stage they are all uncalibrated.

Financial techniques familiar within traditional pricing approaches are essentially reshaped into a framework that lends itself to high-throughput computing (HTC) technology. These matrices can be thought of as cross-sections through the more familiar graphs that show risk factor evolution along paths or lattices; in practice they will be 512×512 or $1,024 \times 1,024$.

The matrix approach allows modellers greater flexibility when specifying the model dynamics, since they are not restricted to a tree/lattice, a path or a closed-form solution.

Step 2: Calibrate risk factors

Select a global set of calibrating instruments with market quoted prices that are affected by each risk factor in isolation:

- Interest rates – swaptions, caps and floors, flow and callable constant maturity swap spreads.
- Forex – European-style and barrier options.
- Equity – American-style options, variance swaps, credit default swap spreads.

Run an optimisation algorithm to simultaneously match the observable

prices with the simulated prices and to minimise the difference, updating the model parameters and probability matrices at each step. Iterate this process until the revalued calibration instruments match the market prices as closely as possible.

At the end of this cycle, the probability transition values described in Step 1 are globally calibrated – these are the ‘probability kernels’. These matrices contain the necessary information to generate future scenarios, similar to the practices of many risk systems, except the estimations derive from mark-to-market values of derivatives, as opposed to historical time-series data. The instruments used for calibration will all reproduce their mark-to-market values when priced off these future scenarios.

Step 3: Initial valuation

Positions and portfolios can now be valued on graphical processor units (GPUs) using backward induction if using single factor, or central processing units (CPUs) using Monte Carlo for multi-factor.

The universe of possible portfolio position values for every position and portfolio at every time point and each risk factor state is now known – this cube of data is the bank’s risk landscape.

Step 4: Orchestration

The first three steps all demand centralised GPU-based HTC. These may be performed at a set frequency – for example, hourly or daily – and can then be distributed to users of this data, who then process it using traditional high-performance computing technology (CPU/grid).

Monte Carlo simulation of portfolio revaluations can be performed quickly by averaging the values attained out of the pre-computed data from Step 3. Similarly, it is possible to extract sensitivities (Greeks), value-at-risk and exposures (potential future exposure and loss given default).

product introduction becomes the question of whether a risk factor is properly modelled or not. But should a bank be trading a risk factor that it doesn’t understand?

■ *Regulatory approval.* The regulatory framework for advanced modelling, where each financial product is independently certified, makes the current problem systemic. By taking a product-centric approach, regulators are cementing modelling inconsistency into each bank. But assuming regulators will accept a well-tested and proven global valuation framework, once this is validated by a core set of models, it is reasonable to assume the ongoing effort in new product approval will be significantly reduced. It may be possible to introduce new strains of products without formal regulatory sign-off, and still benefit from efficient use of capital since they are leveraging the same risk factor models.

Although the above justifications are very good reasons for pursuing a global

valuation framework, there will be substantial hurdles to realisation. People are generally resistant to change, but in this case the stakes are unusually high.

All financial organisations are wedded to the current pricing practices – they are well accepted and there has been a huge investment in their implementation. Everybody has a vested interest in this issue and careers have been built around current modelling techniques. It is difficult for those who have the right level of experience to discuss these topics objectively. However, this transition provides more opportunity than threat – it fundamentally increases the possibilities for banking.

Restructuring valuation technology along these lines makes a lot of sense in its own right, even if HTC is not being considered. But as HTC technology hits the mainstream, there is the opportunity to achieve a lot more than just a simple speeding up of existing pricing technology. If you see HTC as a turbo-charger, rather than as an entirely different type

of engine, then you’ve missed the point. If you’re looking for HTC to provide incremental performance gains, then you’re underplaying your hand.

The application of HTC/GPU technology can profoundly affect financial products – it is potentially the biggest game-changer since Black-Scholes. It is an opportunity for investment banks to create a first-mover advantage that can propel them for decades.

There is a substantial threshold to be crossed to realise global valuation. It is a significant transition for organisations to undertake, and demands bold leadership. But the vision is compelling and the opportunity presents an exhilarating prospect for our future. ■

Claudio Albanese is visiting professor of mathematical finance at King’s College London, Guillaume Gimonet is head of research and development within quantitative strategies at Credit Suisse and Steve White is chief executive of RiskCare. Email: steve.white@riskcare.com