

An Introduction to Global Valuation

Claudio Albanese and Steve White present the possibilities that High Throughput Computing (GPU Technology) opens up for risk management, and discuss the implications for investment banks in the future.

Even when your pricing models are right, they're wrong

Every bank uses pricing models that are specific to the instrument that they are valuing, but even when they are perfectly implemented - perfect data and properly utilised – on a firm-wide basis there are inconsistencies between the models used for each instrument.

Every valuation formula embodies two generic weaknesses;

either explicitly or implicitly, they make assumptions about the future states of underlying risk factors, and therefore the modelling of the future is specific to each individual instrument, rather than specific to the risk factor, and it is therefore not standardised across the whole firm

they are calibrated to a set of chosen observable market prices, however the number of instruments used for calibration is limited and instrument-specific, rather than more broadly representative of the market

So while each position is “correctly” calibrated in isolation and is therefore “correctly priced”, from the broader firm-wide perspective, this instrument-centric approach to modelling bakes inconsistencies into the valuation process impacting hedging activities, the balance sheet and capital adequacy provisions.

This is not “model risk” in the traditional sense - incorrect specification, usage or application - but in a much deeper sense; 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? Does inconsistency necessarily mean that 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 Theory of Finance, prices obtained by consistently using a single model are guaranteed to be arbitrage free, and 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 it, and if we do, can we implement it?

Such a model has not been computationally possible hitherto, but with the advent of High Throughput Computing (HTC – the generalised application of Graphical Processor Units), the amount of processing power available for mainstream industrial applications at commercially attractive rates is dramatically increased.

But GPUs do not work like CPUs, and so the process 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 computationally impossible so far. Many classical mathematics tools such as special functions theory, group theory and spectral theory have been developed to work around this problem, and these techniques lead to Mathematical Finance as we know it today.

The reliance on analytical tractability imposes a straightjacket on modelling possibilities, limiting the flexibility and comprehensiveness of models, and their ability to achieve economically realistic representations. The original justification for these techniques is that they could be realised using widely available technology, however with HTC that justification is diminished.

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.

The good news is that an HTC-friendly solution does exist, solving both our problems above – consistent modelling of future states and broader calibration – and can be achieved by a restructuring of our valuation process flows.

To do this we need to decompose our existing pricing models 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 (books, counterparties...),

- post-processing these results to determine the 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 that this practice is nearly universal is a consequence of the evolution of the industry.

Not only does this decomposed approach enable centralised and therefore globalised risk factor modelling, it also provides the spring board 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.

Rather than calibrating specific models for each instrument, we are now calibrating the risk factors upon which they are priced. But how is this realised?

CHANGING FROM INSTRUMENT-SPECIFIC TO GLOBAL VALUATION

The diagram below depicts the generic change in process flows between the two approaches; at the highest level, all valuation systems are the same – a set of inputs are fed into a pricing function which 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.

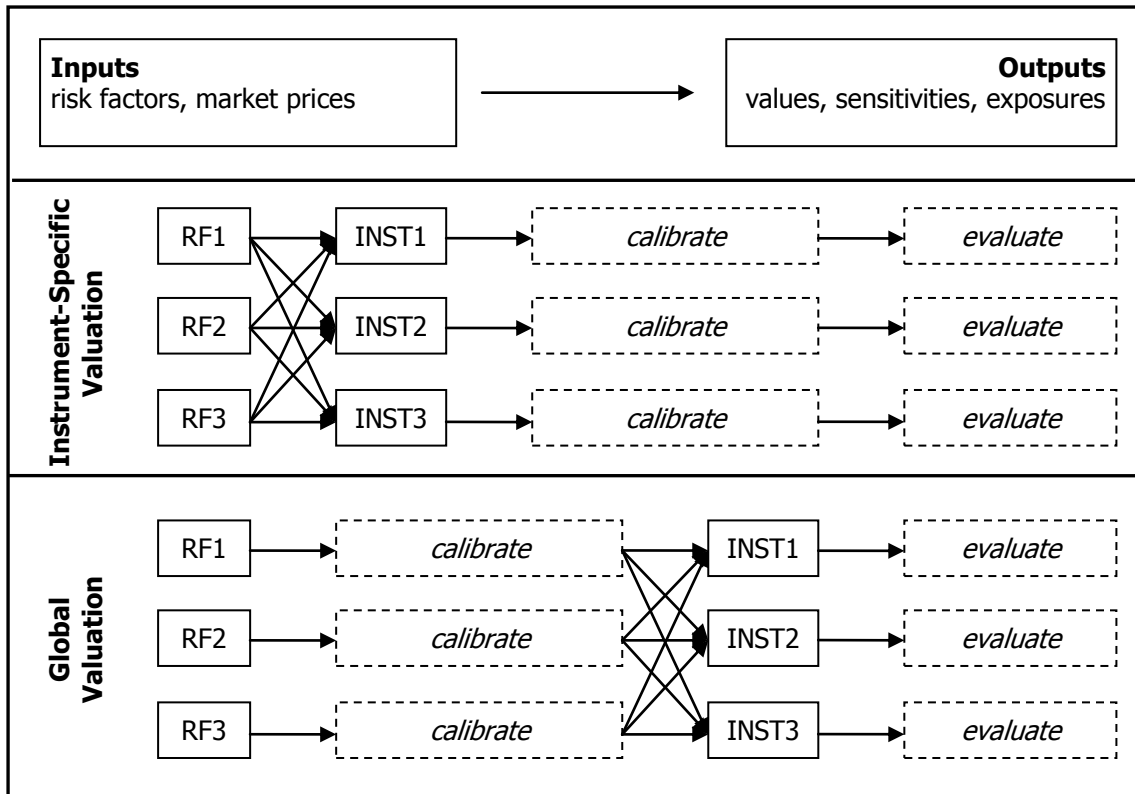


Fig1 – Changes in Process Flows between Instrument-Specific and Global Valuation Techniques

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)

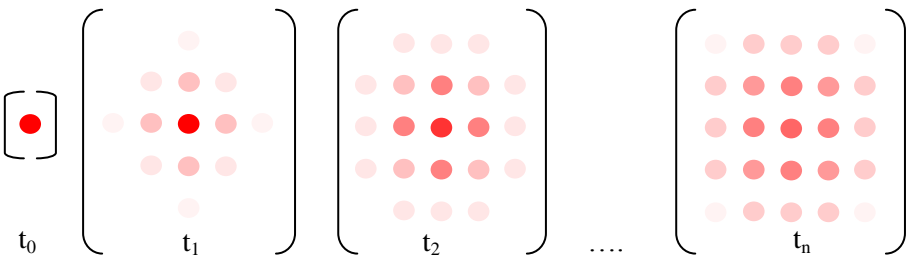
The bulk of the functionality in the pricing algorithm is removed from the instrument-specific valuation such that only the payoff logic remains. Instead, the majority of the analytical and computational effort goes into simulating each risk factor. A 'risk landscape' is created where every possible future state of every risk factor is computed and is probability weighted, and then every position is parsed across it.

THIS IS HOW YOU IMPLEMENT GLOBAL VALUATION

There is no single "right" way of implementing a global valuation model – the framework described here is intended to meet our modelling objectives, and to do so by defining the problem in such a way that it can leverage the computing capabilities of HTC.

This description of the methodology 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.

Step 1 – Risk Factors Initial Setup



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 25bps.

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

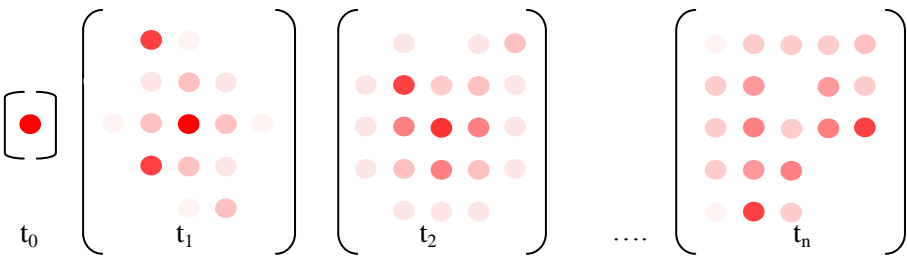
A transition probability matrix is generated based on the initial risk factor process parameter estimation for each epoch. Transition matrices for each epoch are generated by multiplying the single step transition matrix.

The result is a collection of 3d matrices - one per risk factor – but at this stage, they are all un-calibrated.

This illustrates that we are re-shaping the same financial techniques we're familiar within in traditional pricing approaches into a framework that lends itself to HTC technology. These matrices are cross-sections through the graphs showing the risk factor evolution that we are more familiar with; typically they will be 512x512 or 1024x1024 in practice.

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

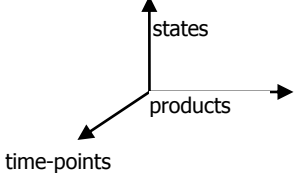


Select a 'global' set of calibrating instruments with market quoted prices that are affected by each risk factor in isolation.

- IR – Swaptions, caps and floors, flow and callable CMS spreads
- FX – European and Barrier Options
- Equity – American options, variance swaps, CDS 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 re-valued 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 our 'probability kernels'. These matrices contain the necessary information to generate future scenarios, similar to the practices of many risk systems, except that 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.

| | |
|--|--|
| <p>Step 3 – Initial valuation</p>  | <p>Positions and portfolios can now be valued on GPUs using backward induction if single factor, or CPUs using Monte Carlo if multi-factor.</p> <p>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.</p> |
|--|--|

| |
|--|
| <p>Step 4 – Orchestration</p> <p>The first three steps all demand centralised GPU-based High Throughput Computing – these may be performed at a set periodicity eg. hourly or daily, and then distributed to users of this data, who then process it using traditional High Performance Computing technology (CPU / grid).</p> <p>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 risks and exposures (PFE, LGD...).</p> |
|--|

GOING GLOBAL - WHAT YOU STAND TO LOSE....

Instrument-specific valuation – the way the industry works at present – has come about with good reasoning 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 (simply 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.

All these facts justify why we do what we currently do, but we need to be prepared to bin all this if we are to adopt the approach described above.... is it worth it?

GLOBAL VALUATION – THE OPPORTUNIY KNOCKS

Since HTC computing in effect means starting over again, there need to be some very good reasons for doing so. Not only are we throwing stuff out, but we are taking a lot more on.

So what are the business benefits that we can expect from global valuation?

1. For Senior Management
 - 1.1 A Second View

An economically consistent model offers an alternative view point for valuation and arbitrage checking – valuations obtained from a global valuation system can be compared to those

obtained using existing systems. It's unlikely that these discrepancies will identify profitable trading strategies, since 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, very valuable information; irreconcilable data ought to give advance warning to senior management of looming problems within their business.

1.2 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 use of their balance sheet, hedge out undesirable risks and can apportion capital to accentuate those risk factors / products / counterparties that are able to achieve the highest risk-weighted returns.

1.3 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 P&L – 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.

2. In The Front Office

2.1 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 through standard APIs with no detriment to computational performance.

2.2 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, VaR-based margining, intra-day full revaluation, real time risk decomposition....

2.3 Better 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.

3. In Risk Management

3.1 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 global calibration effects come into play.

3.2 Integrating Front Office Valuation, Market and Credit Risk

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

4. Better Economies of Scale

4.1 Organisational Effects

Taking a holistic, firm-wide approach to valuation, as opposed to an instrument-by-instrument approach, actually simplifies the challenge when considered in its entirety. The sophistication and volume of the maths is significantly reduced. Instead, we have increased the demands made upon technology and management, but this is centralised across products and functions throughout the organisation; front and middle office can use the same system, trading and risk can use the same system, rates and equities can use the same system, flow and exotics can use the same system, senior management and individual desks can use the same system.

4.2 More Efficient Coding

True polymorphism can be attained by maintaining a single computational engine with core functionality that is shared across product lines and asset classes. By relaxing the constraint of analytic solvability, the technical differences between rates, FX, equity and commodity modelling are reduced; commonalities can be shared in base classes, and the differences can be implemented in derived classes

GLOBAL VALUATION IN PRACTICE – 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 we can anticipate some of the likely organisational impact;

i) The Realisable Power of Executive Management

Executives will get more control since risk is more visible and their organisation is more adaptable, working not just with control variables of products and sectors, but can overlay these with bets on the risk factors that a bank wants exposure to.

ii) The Rise of Risk

The most significant aspect of organisational change that global valuation presents is that risk factor modelling necessarily becomes a centralised activity, and therefore it ought to be owned by the CRO. Quantifying risk is one of the core competencies of an investment bank, and therefore it ought to be value-adding.

The CRO moves from being a controller to a driver, a creator of value rather than a business overhead. This profoundly changes our motivations for investing in risk competency and technology. Risk deserves to be invested in for profitability to the same extent as front office product development, rather than treated as a cost to be tightly controlled.

iii) 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. It will take IT managerial practices into new waters – building bespoke systems will require world class industrial scale management competencies similar to the aviation, energy and pharmaceutical sectors.

iv) New Products Group

Introducing new financial products becomes a significantly reduced effort – it is simply a case of ensuring that the payoff description is accurate, and then applying this to pre-existing risk factors. It is no longer the case that a full revaluation model needs to be developed, distributed and validated – modelling and testing the payoff is a smaller effort.

The bottleneck to product introduction becomes the question as to whether a risk factor is properly modelled or not. But should a bank be trading a risk factor that it doesn't understand?

v) 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 that 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 that 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.

vi) Front Office Quants and Technology Vendors

Although the above justifications point to some 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, there has been a huge investment in their implementation. Everybody involved has a vested interest in this issue – careers have been built around current modelling techniques – it is difficult for those persons who have the ability to discuss these topics to do so dispassionately.

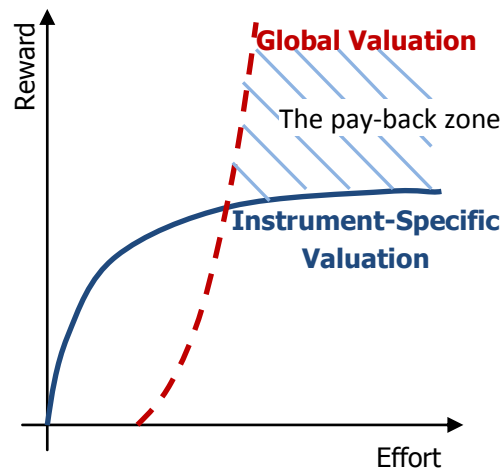
However, this transition provides more opportunity than threat – it fundamentally increases the possibilities for banking, and its early days to see the directions it will lead us in.

The Rewards and Risks of Global Valuation

As shown, restructuring your valuation technology along the lines described above makes a lot of sense in its own right, even if HTC is not being considered – this is a 'right-structuring' of existing practices.

But as HTC technology hits the mainstream we have the opportunity to achieve a lot more than just a simple speed 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 combination of a global valuation framework as described above and HTC is a disruptive technology in the making.

All in all, the application of HTC / GPU technology has the ability to profoundly affect financial products – potentially the biggest game-changer since Black-Scholes. It's 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's a significant transition for organisations to undertake and demands bold leadership. But the vision is compelling; this is how HTC will impact our industry, this is how we will have to respond, the upside is unlimited, and the opportunity presents an exhilarating prospect for our future.

Steve White is chief executive of RiskCare, a consultancy in financial technology. Email: steve.white@riskcare.com.

Claudio Albanese is Visiting Professor of Mathematical Finance at King's College London – for more information and technical papers visit www.albanese.co.uk