How Do Dealer Banks Price Derivative Products?

By understanding dealers' pricing of derivatives, a corporate treasurer can get a better grasp on the true counterparty risks underlying these complex financial instruments.

by Scott Sobolewski

It is no secret that financial intermediaries are still struggling to adapt their business models to the heightened regulatory expectations stemming from the 2008 credit crisis. Banks in the United States have invested heavily in compliance and risk functions to meet time-sensitive Dodd Frank and Basel III deliverables that have garnered much media attention over the past several years. The annual stress-testing requirements for the nation's largest banks (comprehensive capital analysis and review, or CCAR) and the restriction of proprietary trading activities (Volcker Rule) have been particularly impactful.

Taken together, these and other regulations have been reducing banks' margins. In some cases, the regulatory environment has even led financial institutions to suspend previously profitable activities, due to both increased compliance costs in the form of staff and reporting, and increased return hurdles from higher capital requirements. Banks have also been evolving independently of new regulations and have learned from previous mistakes in their pricing of credit risk, funding costs, margin requirements, and capital—look no further than the failures of Bear Stearns and Lehman Brothers in the United States.

Thus, dealers have begun to factor a plethora of bilateral valuation adjustments into every derivative quote. In this new reality, the cost of doing business with large dealer banks is permanently higher, and nowhere else will this cost be felt more than by non-bank financial companies and corporate counterparties trading or hedging in the non-cleared over-the-counter (OTC) derivatives market.

Why Derivatives Pricing Is in Flux

The OTC derivatives market has already changed substantially as a result of several key regulatory reforms. Most relevant at the moment is the Basel Committee's bilateral initial margin requirements on derivatives that are not centrally cleared (BCBS 261).

Starting on September 1 of this year, the largest U.S. dealers (as defined on page 11 of this PDF) and their end users will face margin requirements upon initiating a derivative trade. The size of this initial margin – which, despite its name, is recalculated and adjusted at least every two weeks – will depend on a value-at-risk measure calibrated to a long price history that contains at least one year of stressed prices. Then these dealers and end users will need to post variation margin on an ongoing basis over the life of the trade, based on the trade's mark-to-market value after the trade's inception. Frequency will depend on details of the collateral agreement, but will generally be daily for trades with the large dealers.

Now, banks and end users have been exchanging variation margin for years, and calculating the amount of initial margin at trade inception is not particularly difficult, since organizations like International Swaps and Derivatives Association (ISDA) and AcadiaSoft have taken leading roles in standardizing a single, undisputed view of initial margin for each counterparty. In contrast to variation margin, both parties to a derivatives contract have to post initial margin, and the initial margin received cannot be re-used as collateral with other counterparties. It therefore introduces a cost that cannot be evaded or netted away. On the other hand, initial margin reduces credit risk, as it functions as over-collateralization beyond the outstanding current exposure (which is covered by the variation margin), and it therefore reduces the credit risk (expressed by CVA) and capital cost (represented by KVA). It is therefore even more important to understand the future dynamics of initial margin and its effects on other cost components.

The tricky part comes in projecting how that initial margin amount will change dynamically over the lifetime of a bank's portfolio, because the amount of initial margin required for a specific derivative trade will vary over the life of the trade based on changes in market conditions. For capital management purposes under Basel III, banks will be required to project what their potential future exposure is, then determine how a trade's initial margin is mitigating that exposure at each point in the future. These calculations present an array of mathematical and computational challenges unique to each dealer bank.

Additional challenges for banks that deal in derivatives come from the fact that the Basel III regulation restricts eligible forms of collateral to highly liquid assets, segregates initial margin, and enhances collateral documentation requirements (credit support annexes, or CSAs). The operational and compliance costs of the Basel III rules are challenging in isolation. When combined with the higher risk and capital standards imposed by regulators, the collateral restrictions are forcing banks to make key investment decisions about where they prioritize quantitative research and development—their limited "quant" resources are now being pulled in yet another direction. As a direct result, some dealer banks have already begun exiting capital-intensive derivative offerings, and some derivative end users have been forced to use imprecise hedges in the less-expensive central-clearing market.

Due to the staggered nature of domestic and international implementation timelines, many dealer banks have not yet settled on the optimal solution to a problem with so many moving parts. For example, European regulators have already delayed the implementation of BCBS 261 beyond the original September 1, 2016, deadline, leaving U.S. banks at a competitive disadvantage until the initial margin rules are aligned. Over the short- to medium-term, banks will continue to innovate in terms of their pricing and risk management techniques in response to recent regulation, and derivative prices quoted by dealers will continue to fluctuate as a result.

Amid this market uncertainty, one common denominator becomes clear: All financial market participants benefit from an emphasis on transparency. Dealer banks require regulatory approval for modeling techniques used in derivatives pricing and risk management, and they will soon be forced to publicly disclose additional detail about on- and off-balance-sheet risks. Additional transparency will help them communicate more effectively with both regulators and investors. The benefits for derivative end users are even more obvious; transparency will give them a better understanding of the methodologies used to calculate the valuation adjustments and margin requirements embedded in their next OTC derivative quote.

4 Factors That Impact Pricing

Before the financial crisis, large dealer banks generally used "black box" models to develop derivative quotes. This is no longer an acceptable mind-set. It's important for end users to understand how banks calculate margin and valuation adjustments (collectively known as "XVAs" due to the range of potential adjustments for credit, funding, capital, etc.). These numbers affect the fair value of every derivative trade, and they vary based on both the bank's and end user's creditworthiness, collateral choices, funding costs, and regulatory capital requirements.

Banks generally report each of the four most prominent valuation adjustments at the counterparty level, across their entire portfolio of trades, to capture the offsetting effects of netting. They calculate these valuation adjustments using a Monte Carlo framework. They simulate a large number of risk factors that affect the price of trades out into the future, revalue the trades at various future points according to the market environment in that future state, calculate the net present value (NPV) of the trade by discounting those future prices according to an appropriate yield curve, then calculate specific adjustments based on the forecast information across many thousands of simulations and trades.

The four key valuation adjustments are:

1. ***Credit value adjustment (CVA).*** Credit risk will never be entirely mitigated in derivatives transactions because continuous collateral monitoring across a large portfolio is operationally difficult. In other words, it's rare for one counterparty to post margin and the other to receive it at the exact time an exposure arises, and it's unrealistic to expect collateral to be exchanged continuously—including minimum/asymmetric thresholds or operational problems/disputes—over the life of a derivative. The CVA is the amount by which the actual derivative price, adjusted for credit risk, deviates from the price of the ideal, perfectly collateralized, continuously margined derivative. Under IFRS 13, CVA must be reflected in derivatives' valuations for accounting purposes, not only for banks, but also for corporates, insurance companies, and pension funds that use derivatives.

2. ***Funding value adjustment (FVA).*** All derivative transactions bear a funding cost/benefit for collateral posted/received. In a collateralized trade, a dealer with negative mark-to-market must borrow funds at its unsecured borrowing rate to post collateral to its counterparty in the trade. It will receive interest from the counterparty at the rate specified in the CSA, which is typically the relevant overnight indexed swap (OIS) rate. If, instead, the mark-to-market is positive, the dealer will receive collateral and pay interest at the CSA rate. Since variation margin can be

rehypothecated, a funding benefit ensues.

FVA is a measure of the expected funding cost over the life of the trade. FVA essentially measures the asymmetry between the dealer's unique funding cost—so the price the dealer pays to obtain collateral it will post—and the common rate specified by the CSA, which is what it receives in exchange for posting that collateral. The difference between these rates is going to figure into the dealer's pricing of the trade. FVA is in particular a factor in the pricing of trades in which no collateral is exchanged. Because the dealer will likely need to enter into an offsetting hedge to manage the risks of the trade, and that hedge will likely be collateralized, the dealer may allocate to the original, uncollateralized transaction the funding asymmetry between the hedge and the price quoted.

FVA is not yet an accounting requirement, though many of the largest global dealers consider it a true contribution to their derivatives' fair-market values and publicly disclose its value alongside that of CVA.

3. ***Capital value adjustment (KVA).*** Some dealers price their unique cost of regulatory capital into each derivative quote. Generally the KVA metric captures three components: counterparty credit risk (CCR) capital, CVA risk capital, and market risk capital. While the CCR capital charge and CVA capital charge can be computed along netting sets and then aggregated, the market risk charge depends on the entire bank's portfolio. In order to calculate the future capital requirement of a derivatives portfolio, it is necessary to project all three types of capital charges over the entire life of the current portfolio. This makes KVA the most computationally intensive and least widely used valuation adjustment at the moment, though it is gaining support within the dealer community as banks continue to report and allocate regulatory capital more frequently and efficiently. However, KVA has not yet found its way into banks' financial statements.

4. ***Margin value adjustment (MVA).*** Driven by Basel's bilateral initial-margin requirements for derivatives that are not centrally cleared, as well as the general regulatory push toward central clearing, MVA is similar to FVA, but MVA is applied exclusively to the asymmetry in funding costs for posting initial margin, whereas FVA applies to variation margin. Like FVA, MVA values are driven by the difference between the dealer's unsecured funding cost and the interest rate received on margin posted. The primary difference is that MVA is always a cost without the potential of netting effects. However, MVA is much harder to determine because initial margin is, in itself, dependent on future projections of mark-to-market exposure. (While FVA is dependent on a Monte Carlo simulation, MVA requires, in theory, a Monte Carlo within a Monte Carlo, which is computationally very difficult to accomplish.)

In calculating these four valuation adjustments, banks require risk factor evolution models specific to each individual risk factor; categories of risk factors that affect the price of the typical trade include interest rates, FX rates, consumer price indices and real rates, default and recovery rates, equity prices, and commodity prices. The banks also take into consideration the volatility of each of these prices. The portfolio nature of the value adjustments also makes it necessary that all risk factors within one portfolio are simulated simultaneously in a fashion that reflects their correlation correctly. In addition to the scenario evolution models, banks also require product-specific pricing models from which they calculate the mark-to-market at each future date. Ultimately, this framework provides a set of NPVs for the portfolio that reflect trades by future dates, under different market scenarios. The bank can then base valuation adjustment calculations on these NPVs, with varying degrees of statistical confidence.

Simulating these four valuation adjustments for a portfolio of trades, over thousands of random or semi-random market forecasts, becomes computationally expensive very quickly. To give some perspective, typical dimensions could include 10,000 trades, 120 evaluation dates (quarterly time steps over 30 years), and 10,000 market scenarios. That data cube alone would contain 12 billion NPVs! If pricing a single trade took an average of 50 microseconds, the calculations in our example would require around 170 CPU hours on a single-core machine or two and a half hours on 64 cores. This does not yet include the extra cost of computing sensitivities either within the exposure simulation (as required for MVA or KVA calculations) or of the end result (as required by a desk trying to hedge the XVA risk). Dealer banks generally strive to produce daily or weekly risk figures for senior management review and for hedging purposes. Thus, dealer banks are putting a premium on calculation speed and model performance.

Many of the smartest minds on Wall Street are currently working on ways to shave microseconds off calculation times using a combination of better hardware and analytics methodologies. There is some room for optimism among end users, as advances in hardware (graphics processing units) and analytics methods ("short cut" approaches that avoid time-consuming "brute force" activities) have the potential to lower the cost to dealers of derivatives trading, which would encourage more competitive pricing from the efficient dealers.

Additional Considerations

Once a bank has calculated the data cube of projected NPVs for its derivatives portfolio, the bank is ready to calculate its valuation adjustments for a single trade or counterparty. In doing so, it must be careful to take into account any unique features of its individual collateral agreements, such as interest rate floors, minimum transfer amounts, one-sided thresholds, margin-call frequency, and optionality for collateral type or currency. Concentrations in credit or funding have the potential to affect trade pricing for specific counterparties due to netting, and unique collateral arrangements make it more difficult for dealers to hedge offsetting risks with other counterparties.

As a result, dealer banks have been pushing to standardize CSAs, in an effort to more effectively hedge their portfolios. They are often willing to pay a large premium for removal of certain old-style derivatives features currently considered "in the money." For example, some pre-financial-crisis CSAs have interest rate floors on cash collateral posted to either counterparty. The floor value is usually set to zero, so when interest rates are positive, the feature is practically irrelevant; counterparties collect interest on any cash they post to the collateral recipient.

However, in negative interest rate environments, this feature prevents a counterparty from paying interest to the collateral recipient on cash that's posted as collateral. For derivative traders that have large negative mark-to-market positions in Europe's current negative interest rate environment, this feature could be immensely valuable. If those traders' counterparties sought to renegotiate a CSA to remove its interest rate floor, they would have to pay a premium for its removal to compensate the trader for the fact that it will now need to pay interest on cash collateral it posts over the remaining life of trades governed by that CSA. The value of this feature is calculated as the difference in portfolio-level FVA between all trades with the interest rate floor and all those without it, demonstrating the power in knowing how such valuation adjustments are calculated.

Although the majority of affected banks have already implemented Dodd-Frank here in the United States, many are still exploring the most efficient ways to respond to regulations and persistent demands for increased transparency. The choices dealers make will have profound effects on their ability to effectively communicate with regulators, investors, and derivatives clients.

It has become clear that a robust understanding of pricing and risk on both sides of a trade will go a long way toward keeping markets liquid and vibrant. To that end, Quaternion Risk Management will release an open-source version of our Monte Carlo framework in the third quarter of 2016, with the goal of jump-starting a global discussion on the benefits of derivatives pricing transparency.

The financial concepts underlying derivatives pricing are difficult. But when a wide swath of financial market participants understand how banks are calculating their prices and projecting their risks—and so understand the potential risks of complex derivatives products—that knowledge will help reduce systemic risk in the OTC derivatives market as well as reduce the cost of regulatory compliance. It will also increase liquidity for capital-intensive products. As a broader pool of organizations have tools to efficiently price and manage the risks inherent in these complicated products, more financial market participants will feel comfortable trading them with large dealers. Liquidity will increase in the OTC derivatives market, and increasing numbers of corporate treasurers will be able to hedge financial risks more precisely than they could in the central-clearing alternative.

Scott Sobolewski *is a principal consultant at Quaternion Risk Management and splits time between corporate offices in New York and Boston. He specializes in capital planning, stress testing, and model development at large U.S. banks, and advises financial institutions on risk management and regulatory compliance matters.*

For More Depth on This Topic…

In the spirit of full transparency, the founding partners of Quaternion Risk Management recently published a comprehensive guidebook to pricing and risk for derivatives and structured products in the modern, post-crisis regulatory environment. As ex-bankers and quant risk managers, the authors of *Modern Derivatives Pricing and Credit Exposure Analysis* provide a detailed explanation of mathematical theory and practical approaches that drive pricing and risk across every asset class—interest rates, foreign exchange, inflation, credit, equities, and commodities—allowing for a functional understanding of pricing and risk on both sides of a potential trade.