SOCIETY OF ACTUARIES

Article from:

Product Matters!

October 2010 – Issue 78

SOCIETY OF ACTUARIES Product Development Section

Product Matters!

ISSUE 78 | OCTOBER 2010

Market-Consistent Pricing
As the Market (Sort of) Normalizes
Separating the permanent from the temporary grayness
Part 1 of 2

By Eric Clapprood and Mitch Katcher

The influence of market-consistent pricing proponents was growing with perhaps more momentum than ever in 2008, when the bottom suddenly fell out of the financial infrastructure of the U.S. and global economy. "*The resulting freeze in credit markets and lack of transactions in previously liquid instruments showed that just when the 'tail' we all feared emerged, the data became unavailable to calibrate to, rendering market consistency meaningless."*

Wait—hold on. "*Actually, those who had transacted prior to the crash and covered or transferred risk off their balance sheets embodied the principles of market consistency with those very transactions, and the benefits of those decisions clearly show that market consistency is the only way to price."*

No—just a minute. *"Now that we've seen this crisis play out for a year-and-a-half, we know that there were incredibly volatile and irrational moments during that time that should not have been reflected in valuations because they were simply not credible, demonstrating the flaws of market consistency for a solvent insurer whose view is long-term, not day-to-day."*

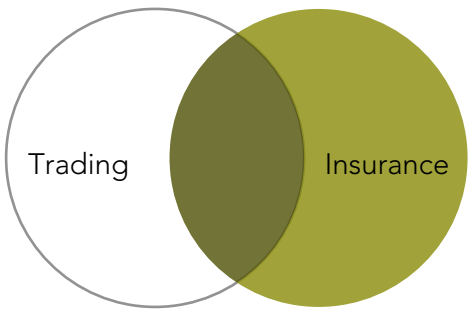

CONTINUED ON PAGE 5

This is the debate occurring in that gray area where the black insurance industry circle overlaps with the white market consistency circle. Indeed, those colors are precisely how some market purists would describe the convergence of insurance and trading (after all, there is no market without trades): Trading is clear and transparent while insurance is a black box.

GrayPixels

The gray area where insurance and trading are overlapping consists of many pixels, but unlike those in your flat-screen T.V., not all pixels are the same size here. The larger ones are:

- Long-dated (and other obscure) measurements
- Regulatory requirements and accounting differences
- Short-term volatility
- The non-equivalency of traders

Some of the above pixels will clarify over the next two to five years, as market-consistent pricing "technology" rolls out its new T.V.s, but others will likely remain perpetual challenges: in the end, many insurance products' market-consistent pictures will always be an artist's rendering.

Long-dated (and other obscure) measurements

Principle: *There will always be an area of the consumer market that exceeds the horizon of the traded market.*

VA Guarantees: The neo-classic market-insurance hybrid product class is that of the Variable Annuity (VA) riders, the GMXBs that guarantee a payout despite the performance of underlying (mostly) equity-based separate accounts. Consider the attempts by regulators to appropriately value these instruments. One needs to look no further than the United States to find that conclusions drawn and implemented after years of debate by industry experts, all focused on complex, stochastic-calculations-based answers, end up in completely different territories. The NAIC's answer to the challenge was C3 Phase II capital, followed by AG43 (VACARVM) reserving, both of which (in simplistic terms, ignoring the Standard Scenario) set "tail of real-world distributions" as the definition of valuation. Fairly simultaneously, the FASB came to an entirely different conclusion under FASB's *Accounting Standards Codification,* Topic 815, Derivatives and Hedging(formerly SFAS 133) and then *FASB's Accounting Standards Codification* Topic 820 - Fair Value Measurements and Disclosures (formerly SFAS157): fair value, which uses the "average of a risk-neutral distribution" to send pictures out to investors and the public in general. (The risk-neutral approach is called such because the investor theoretically is only concerned about the average of the distribution, not the tail. While this is not generally reality—investors are considered to be risk averse—the model simplification employed includes adjusting the implied vol input to reflect the appropriate price.) Suffice it to say that the two U.S. approaches are very different, and that the reason for this, more than any other, is the obscurity of long-dated and other measurements and their interpretation within two different constructs: exit price and current capital levels.

The obscurity occurs on both the insurance and the market sides of the hybrid product. On the insurance side is the often-subjective prediction of policyholder behavior, which itself needs to be divided into two categories: an approximation of randomness and an approximation of efficiency. From a randomness standpoint, the VA riders contain some of the same risks that have always driven insurers' products, like mortality and lapsation. From an efficiency standpoint, actual use of the rider (withdrawals) the policy owner is paying for is one of the most sensitive pricing components product actuaries need to tackle. At the center of the market consistency/behavior debate is the issue of efficiency of the option holder. While some argue the option holder is always efficient in an options world, the reality is that if a block of business can be sold with certain inefficiency expectations (e.g., lapses and less than full utilization) then that is, in fact, the market, which is a result of an option being attached to an insurance host contract. Certainly there is evidence to show that in an economic downturn some policyholders will lose their jobs and need to access funds to the extent of a full lapse of a VA contract that was in the money. This policyholder did not look to maximize a Black-Scholes formula, but simply needed to pay the mortgage, and he helped define the market.

On the market side, those who have run hedging programs or dealt in any way with the valuation of long-

CONTINUED ON **PAGE 6**

Eric Clapprood, FSA, CERA, is senior manager for Deloitte Consulting, LLP. He can be contacted at eclapprood@deloitte.com.

Mitch Katcher, FSA, MAAA, is principal for Deloitte Consulting, LLP. He can be contacted at mkatcher@deloitte.com.

dated guarantees will relate to the image of a field scout looking through binoculars at the horizon, having a less and less clear read on implied volatility (vol), until such time that it is clear that there is no line of sight to certain points down the road. What is less often discussed is something so close that the field scout is actually standing in it: the correlation assumptions, beginning not 30 years from now, but today. Such correlations aren't found in newspapers or on Bloomberg screens as easily, but can be traded. One can choose to enter a trade such that you are paid if the rate and equity correlations in the future are higher than X, and you pay the bank if the correlations are lower than X—similar to a futures trade. The level of X is arguably where the market sees correlation in the future. Generally, a higher correlation of rates and equities is bad for the embedded guarantees, and the implied correlation of these trades has been at or above 30 percent for some time. What happens if an insurer is using a model calibrated to historic performance that results in a negative correlation? The set of questions that need to be addressed are similar to those of the long-dated volatility issue.

UL Guarantees: The obscurity issue is not limited to VA measurements. With Universal Life (UL) secondary guarantee products, insurers need a view on where rates will be as far out as 50 or more years in the future, which extends beyond the liquid markets' view. The products will assure that despite poor investment returns, contracts with certain minimum premium payments will stay in force. There's a risk that if rates quickly rise at a certain point, bond (fair value) prices collapse, and there's a risk that long-term rates are too *low* to support required investment returns. What seems simple at first becomes quite a complex series of interest rate puts and calls.

Therefore, in addition to a view on expected rates beyond the typical horizon of the market's binoculars (30-year bond issues and 40 years of increasingly-illiquid futures trading), and perhaps more importantly, is a required view on the volatility of those rates.

For the last several years, insurers have been (mostly) quietly experimenting with what fair value would look like on these long-term rate options. A typical approach to a stochastic rate generator is a two-factor Hull-White model that calibrates to market prices. One need not understand the specifics of such calibration to appreciate the phenomenon that has bothered life actuaries but been overshadowed by the VA challenge. (Be there no doubt that the overshadowing is largely due to the mark-to-market of VA riders under GAAP and a lack thereof for life insurance rate guarantees.) In particular, when one calibrates the major variables of a stochastic rate generator (the volatility and mean reversion factors) to observable rate options (typically five- to 10-year swaptions) and runs a model 40 to 50 years or longer, the rate set will be very high. *Average* rates after several decades will look more like *historically high* rates.

This is because the mean reversion needed to counter the high volatility of rates becomes stronger and stronger as the projection period increases. Part of the reason for this is that while rates are "allowed" to go negative there remains a sensible and economically-explainable bias toward positive rates, skewing returns more in the upward direction than in the downward one, to meet the implied vol requirements associated with traded swaptions.

But the real reason for this is that there is a lack of calibration data on the long end of the curve. If the market were readable for 40-year and 50-year caps and floors, there would be more calibration points that most definitely would revert rates back to "normal." This would not fully solve the problem, though. Using the 40- and 50-year options data (if it existed) would be fine for 40- and 50-year views, but it would then wreak havoc on the shorter end of the curve. The only solution to calibrating to what one thinks 40- and 50-year options might trade for and what five- and 10-year options are trading for is to make the models more complex than they are today—by quite a lot.

In other words, to create market-consistent long-dated rate guarantee pricing one has the challenge of (1) not having enough data for calibration, (2) finding perplexing and completely unreasonable results when calibrating to the data that does exist, and (3) considering the daunting task of increasing the complexity of models exponentially if one day the correct calibration exists.

The temporary grayness here, in our view, relates to a number of methodologies employed by insurers issuing economically similar guarantees, but interpreting the "observable data" differently, often because the data each has available to them is different.

The most visible gnat here is probably implied volatility assumptions for long-dated embedded derivatives. An insurer not trading in long-dated options may not observe long-dated volatility, whereas one making those trades will. In the absence of a liquid, observable market, historic volatility might be used. Some insurers may combine market data with historic returns. A number of questions remain unresolved.

If options are attached to non-tradable host contracts, how does one translate un-attached options of similar construct to the option being valued? One might argue that a pure translation must be made from observable volatility levels for a similar-term option. One might also argue that there is a lack of direct sight into the options market, which is related to, but different from, the "embedded options market."

Some insurers have chosen to use historic volatility throughout the implied volatility curve; others have used market implied volatility as long as 15 years into the projection; and some choose a point between five to 10 years through which market data is used. In all cases that include some use of market implied data, the next question becomes what to do after that data becomes no longer observable. Here there are a finite number of choices: grade to historic volatility; hold the last observed spot or forward vol constant; extrapolate the trend; or develop a method that combines these concepts. A key question underlying this step from what one sees as observable to what one cannot observe is a question of relativity: If, for instance, there is no marketplace for 20-year puts, does that argue for a high vol to be assumed (under the presumption that "envisioning" such a marketplace logically infers the extrapolation of what is usually an upward-sloped vol curve)?

Any downward slope of the spot volatility curve brings with it a precarious interpretation of forward vol. For example, if it is presumed that 10-year spot vol is 25 percent because this is where options are trading, but 20-year spot vol is 15 percent because it is deemed to be valued based on historic market performance, the only way the insurer's model will "get to" 15 percent 20-year spot vol is to use a vol so low in years, 11 to 20 of the model, that it has never been historically observed. Thus, it becomes impossible to model both a realistic 10-year period (years 11 to 20) and a total period over 20 years that it also deemed historically calibrated in aggregate. This is one reason some who avoid use of the market-implied data altogether might argue their interpretation appears more consistent than those who need to make this "grade to historic" decision.

One development that may help remove the grayness is a subset of financial regulatory reform discussions, which centers around derivatives being brought onto exchange platforms. Currently, if one is not active in the options market (on the asset side), it is arguably difficult to call it observable.

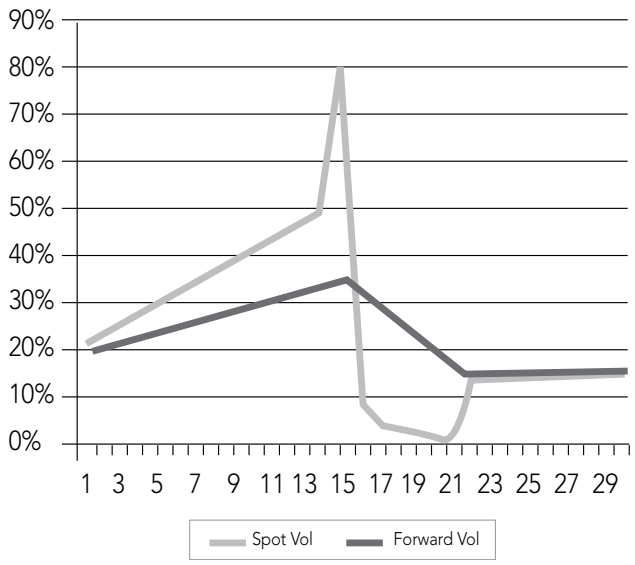

The spot volatility curve shown assumes market volatility rising to 35 percent by year 15 and then grades to historic volatility of 15 percent by year 20. This results in a zero forward vol assumption in some cases, raising important questions for insurers who use such a method.

CONTINUED ON **PAGE 8**

The permanent grayness is the concept of what will likely be the ever-present (albeit ever-changing) set of risks insurers take on that are not liquid and observable with consistency.

On both the life insurance and annuity side, we have products whose components span the spectrum of clearly-observable market inputs to opaque regions of actuarial estimation. Our principle is that this will always be the case. Imagine all components being market-legible, which would be a requirement for a world where this principle is violated. In such a world, by definition, each piece of the product is traded easily, meaning that the product itself is simply a basket of other traded goods, and the life insurance company is acting merely as broker. This is basically the case with mutual funds, a high-growth area for some U.S. insurers. A non-trivial point we'll return to below is the low ROA and high ROE on those mutual funds, and what it says about insurers' choices in a fully fair-value world, which brings us to the topic of regulatory requirements and accounting differences.

Regulatory Requirements and Accounting Differences

Principle: *For the next 20 years it will be impossible for global insurers to avoid regulatory conflicts with market consistency.*

As we noted above, a major conflict exists in just the United States, between the statutory direction and the FASB direction related to the most glaring intersection of insurance and trading, VA rider guarantees. Were this the result of one framework having recently "modernized" its approach and the other not having done so in several decades, a conclusion could be drawn that a chronological gray fog had temporarily descended upon the city of Pricing. Unfortunately, both the NAIC and FASB very recently (especially in terms relative to regulatory change horizons) concluded on their respective guidance after much debate and analysis. Each pronouncement set (C3 Phase II / VACARVM and FAS133 / FAS157) carried with it that fresh-paint smell that comes with the word "stochastic." Both the Monte Carlo simulation required for FAS157 options replication and the multiple servers needed to project principle-based AG43 reserves seem to be making a quantum leap in valuation.

In the end, though, statutory principles are not market consistent, in large part because they were never intended to be. With a focus on liquidity and capitalization, U.S. stat has a different goal. With regard to embedded guarantees in VAs, for instance, statutory reserve and capital requirements prescribe the use of a scenario set calibrated to historic returns, not the market's view of future returns. Insert at this point a debate, if you choose, over whether or not market implied measures (forward rates or implied vols) have ever "done a good job of predicting" the future. If you make this choice, however, be prepared to realize that such accuracy (markets' "predictions" of the future) is irrelevant. The historic return and vol, respectively, of the S&P 500, are around 11 percent and 15 percent, compared to a risk-neutral set that currently will presume (using a 10-year horizon) less than 4 percent and more than 25 percent respectively. Does this mean that the markets are predicting that stocks return, on average, the same as risk-free investments, and that they will experience a standard deviation on average of 25 percent? No. The markets believe that there is a risk aversion that will on average lead to stocks outpacing bonds' returns. The market believes the distribution of stock returns are—unlike the Black-Scholes assumption—not normal. And there are a half dozen other differences between options pricing formulas' assumptions and what the market truly believes. In the end, though, these are all factored into the implied volatility of the options' price, and the average—not the tail—of a normal distribution of returns is used as the value within Black-Scholes. This is an important concept, because when this average is much worse than the tail of a "real world" stochastic set, it says something fairly bold about the difference between statutory and fair value measurements. The issue, therefore, is not whether the market's implied vol is a good prediction of where vol will be in the future, but the fact that it is what the market would use were it to price the liability at hand.

If an insurer sold a 10-year S&P 500 at-the-money put to a policyholder in today's environment (low rates and high vols), it is quite likely that for a notional amount of the put that would sell for $100 at the same moment "on the street," the insurer could sell it for $90, and see the combination of reserves and required capital (even assuming 100 percent efficiency of the policyholder—we are dealing with a true put, not a VA in this example) consistent with AG43 and C3 Phase II be less than the $90 premium, thus generating an instant profit and potentially infinite ROC. If the company hedged the risk by purchasing an offsetting put, the profit and negative strain disappear. Does the disconnect with market consistency remove an incentive to hedge? Put it this way: If one's only consideration was statutory results, then in this example, it becomes difficult to justify the hedge.

There is an exception to the GMXBs' stat-based hedging incentive, which is the subjective nature of determining the long-term cost of a hedging program. Some insurers have assumed that a hedging program should reduce their statutory reserves and capital, while those who have actually modeled such a program find that reserves worsen, and, quite often, even capital worsens, too, when hedging is layered onto a statutory framework.

The accounting differences become further pronounced when comparing to IFRS, and, again, when comparing capital requirements across various regulatory regimes. Consider the differences between RBC in the United States, SMR in Japan and Solvency II in Europe. The combination of IFRS and Solvency II will lead to an essentially market-consistent income statement and balance sheet for European insurers. However, RBC requirements will remain largely factor-based and result in not only differences of magnitude (i.e., in the case of the impact of an equity drop), but direction as well. If one chooses to hedge the fair value balance sheet for equity guarantees and rates increase while vols fall (both "good things" to the market-consistent metrics), then a hedging program that has worked perfectly will result in an asset loss that offsets the liability gain in US GAAP and IFRS. On the U.S. statutory side of the balance sheet, however, the liability is essentially unchanged, while the asset loss carries over, for a net loss that could be significant.

The lack of market-consistent measurements on GMDBs, some lifetime GMWBs (or their components) and UL guarantees is in no small way responsible for the lack of robust hedging programs addressing those risks (compared to those around period-certain GMWBs and GMABs) according to those who help make these decisions. In addition to benefiting from the relative GAAP smoothness that comes with a lack of mark-to-market on such designs (and avoiding cumbersome, expensive and complex hedging processes), insurers have valid concerns borne out by the recent financial crisis regarding capital measurements. In general, if we had to sum up GAAP versus Stat priorities in the United States for insurers with one rule, it would be, "In good times, think primarily about stable earnings; in bad times, protect the capital."

In other words, insurers can legitimately state that there is not one market-consistent metric when there are multiple markets. One can easily envision a simplified scenario whereby, on a market-consistent basis nothing changes for an insurer (due to a well-run hedging and reinsurance program covering liabilities' worsening during a shock), but on a U.S. RBC basis, things worsen dramatically and result in downgrades and even potential regulatory action.

The above situation didn't occur in 2008. In fact, in some ways the opposite occurred. Market-consistent metrics showed much worse damage in most cases than U.S. stat metrics reflected. Hedging gains due to historic rate plunges and vol jumps flowed into statutory results as a buffer. Going forward, however, the opposite risk is greater: that which is outlined above, whereby hedging market consistent metrics will result in losses due to vols falling and rates rising. Most insurers with hedging programs have spent not an insignificant amount of time envisioning their answers to an analyst's question that is basically, "Why has your capital cushion worsened as the market recovered?"

CONTINUED ON **PAGE 10**

The temporary grayness here is, hopefully, the disconnect among accounting principles and capital requirements in the United States, the European Union and Japan. If these three economies' regulators fall in line with each other, much of the rest of the world should follow. In the middle of all of this are the rating agencies, criticized by many post-crisis as being more of a reactionary device than a warning sign. The agencies are in a difficult position, though, regarding taking a view on market consistency versus competing frameworks. If, for instance, an agency believes a fully market-consistent balance sheet is "the way" to go, there remains the risk that local (state) regulators require an action plan or "worse" as a result of traditional RBC ratios falling low (despite no change in a market-consistent balance sheet)—a risk that can't be ignored.

The permanent grayness is probably the inevitable difference of approach among 200 countries. The magnitude of this problem, however, is the least of the "permanent gray" problems we list here, so long as there is an agreement among the Big Three (United States, European Union and Japan).

Up Next: Part 2

In Part 2, we will discuss the phenomenon of short-term volatility; the non-equivalency of traders; and how product development is being impacted by market consistency. □