Risks and Rewards Newsletter

August 2005 – Issue No. 47

Are Your Scenarios on Target?

by David Hopewell

It's summertime, and the living is easy. That is, unless you have been called on to deliver a simulation for product pricing, capital estimation or valuation which is simultaneously understandable, realistic and appropriate to the task at hand. Insurance company products have increasingly become bundles of financial and life contingent guarantees with the expected profit hanging by a thin thread of behavioral expectation. Actuaries are faced with the dilemma of how to incorporate advances in insurance products and modeling technology into their work. Increasingly, this means looking at a large number of potential outcomes using stochastic simulations.

In many cases, determining the profit and risk profile of a product requires estimating the future value of a set of financial guarantees and other incentives jointly with the behavior of a policyholder faced with a complex set of incentives. Policyholder behavior is responsive to the value of various guarantees, but may also include strongly held beliefs, personal needs and advice from the broker or distant relative who sold them the policy. These two issues are individually challenging, but combined create a problem that is more complex than the valuation of mortgage derivatives and structured credit transactions, both of which have attracted massive amounts research and modeling resources. Software has evolved, but in many cases is not keeping up with the creativity of its users. Despite these challenges it is in our enlightened self-interest to minimize any avoidable loss of precision. One area where this plays out is in scenario generation and use.

Simulation problems have a simple schematic that belies the highly detailed nature of what lies beneath. In actuarial models—as in option pricing models–the input is a scenario set. The model is a set of rules that represent the function to be estimated, usually in the form of a cash-flow generator. The metrics from actuarial models can be varied and complex, covering measure of profitability, income, risk and surplus. The metric from an option model (see Figure 1) is a simple average, no matter how complex the underlying function.

turn to page 20

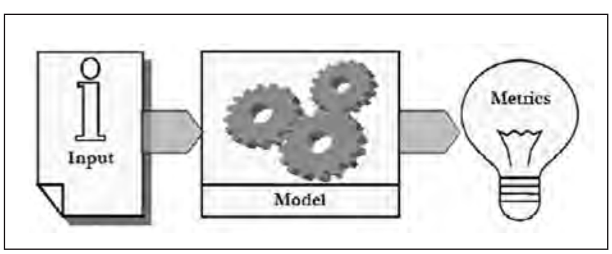

Figure 1

From Page 19

The common decisions for any scenario generation project include the selection of the sample set that serves as the basis for the scenarios, the assumed distribution(s) for parametric approaches or statistical sampling technique for non-parametric approaches and the method of associating marginal distributions. Once those decisions have been made and the process for generating scenarios with those qualities is complete, the final step is to determine parameters that best match the important features of the sample set. That step is called calibration: where theory meets reality.

Risk neutral asset pricing is one of the great advances of modern finance and a common language market participants share.

When Worlds Collide…

Most of us are familiar with the two broad approaches to scenario generation: risk neutral scenarios are used to price options and guarantees, while real world scenarios are used to estimate the potential outcomes or probability distributions of future observable events. Risk neutral asset pricing is one of the great advances of modern finance and a common language market participants share. The great value of risk neutral analysis to the actuarial community is that it provides a theoretically supported shorthand technique for valuing assets by observing the price of other assets. This makes it worthwhile to master, as we develop and often hedge financial guarantees seen nowhere else on earth, but it's important to remain aware of the hazards in determining the "implied" values for uncommon assets from observations of more common assets.

The first step in generating risk neutral scenarios is extracting important distributions implied by observed prices. If we knew what those distribution were and could look deeply enough at the market, our calibrated model would produce scenarios that perfectly reproduce the prices of all observable market instruments, which therefore is highly likely to produce the correct price for a new asset. So why don't we just do that and go home early? The answer is as simple as the solutions are not; models are not comprehensive or accurate enough to reflect reality. Most observed asset prices contain components (usually lumped together and called liquidity premium) that make it difficult to isolate the distributions we are interested in and, worst of all, the implied future distributions generated from different sets of observed asset prices are often inconsistent and mutually exclusive!

An example is the difference between implied forward rate distributions derived from interest rate caps/floors and swaptions. An relatively simple interest rate model such as the Brace-Gatarek-Musiela (BGM) model can easily be calibrated to perfectly reproduce the price of at-the-money caps and floors, but is not likely to do a good job of pricing swaptions, especially if they are out of the money. More complex models reduce but don't eliminate this issue, since it is a feature of the observed market seen through the lens of existing models. The calibration instruments do not contain more information than is needed to set their own price.

The same is true of the Black-Scholes derived, implied volatility model for equity options, which may succeed in matching the observed prices of vanilla options, but also gives different values for more complex structures. The most variation in implied prices is frequently seen in strongly path dependent options, where the value ultimately depends not just on what values the underlying achieves, but when. One way this is dealt with is to increase the size of the sample—the market instruments used—and relax the fit (minimizing the sum of squares between actual and model pricing). But this has limits, too. There may be no instrument that contains the higher order information required to accurately value a path dependent insurance guarantee: the market may not be complete.

Here is where reality trumps theory. In the complete market of theory, prices are unique and profit-seeking traders assure that the same risks trade at the same price, no matter the package. In an incomplete market, prices are not unique. There may be correctly more than one appropriate price for an unobserved asset, such as an insurance financial guarantee, because many available strategies with similar cash flows have different expected returns. There is no roadmap to define the best strategy or whether a strategy tests best depends on the model and calibration used to test it. The problem is circular and solving it is art as much as science.

There are various approaches to the problem, and all have their place in an insurance company setting. One way is to use a more theoretical approach—chosen for tractability. The primary theoretical approach uses the lognormal distribution, with a fixed or time varying volatility for equity and fixed income, and a fixed linear correlation between the two. Calibration is relatively simple if simple instruments are chosen, but is more likely to be unstable since future change in volatility and correlation will depend on the market path.

The volatility and correlation structure observed in the current period may be extended as an assumption for future periods. However, this approach is likely to give optimistic estimates when compared to market consistent calibration methods for out-of-the money options. The latter approach can be extended with volatilities that vary by index or rate level as well as by substituting copulas for correlations, however the complexity increases rapidly.

Another approach is to use an alternative model that more accurately reflects the market-implied distributions of observed instruments with fewer parameters. Unfortunately, few alternative models have a closed form solutions for option prices. The Heston GARCH model is an equity option model with a closed form solution that offers a better current and future period fit than Black-Scholes option implied volatility (Heston and Nandi 2000).

If even more realism is required, i.e., combining dynamic real world policyholder behavior with option valuation, much of the financial theory focused on giving convenient answers under simplifying assumptions is lost, and prices need to be determined from first principles of simulating the best strategy and the cash flows expected from it.

Get Real…

Since most of the theory for pricing options has focused on the risk neutral world, it may come as a surprise that options can also be priced entirely using real world models. Consider a delta hedge on a European option. The Black Scholes value of the option can be matched (in the limit) with a risk neutral simulation. It can also be matched (in the limit) with a (simplified) real-world simulation, where in addition to the option premium and payoff, a delta hedge strategy is simulated. When the real-world payoff of the option and the strategy are combined, the theoretical price is obtained. Each path has precisely the same net cost, which equals the initial price of the option. This replicating strategy approach only works if the real world volatility, dividend and interest rate assumption match, but it proves the point; risk neutral pricing is a convenience used when certain simplifying assumptions are acceptable. As those assumptions become less acceptable, direct simulation of the hedging strategy becomes necessary.

The primary simplifying assumption is that the market is complete and any participant who chooses can be relieved of any risk for the same price as any other participant. Writers of unique path dependent options care about this assumption because the market is not complete with respect to their risks. No model or calibration technique can avoid that and so those seeking the best answers may have to return to the real world. Ironically, in the real world the way is more uncertain. Like a giant X in the middle of an empty map, our target may be apparent, but the path and hazards along the way hide in wait and challenge us to overcome them. At first, the journey is easy and the steps resemble those of the risk neutral approach. There is a sample of outcomes that define the target. In the real world the sample is observable history rather than observable market prices. Here is the first challenge. Which history is appropriate and when?

The Heston GARCH model is an equity option model with a closed form solution that offers a better current and future period fit than Black-Scholes option implied volatility.

There are many facets of markets: countries, time periods, instruments, or in the case of indexes, actively managed samples of instruments. There are also markets that have not survived, where claims on assets if not the assets themselves have been destroyed. Since we are generally looking at outcomes that are relevant to our management, shareholders and regulators, it is appropriate to ignore markets that don't survive and keep in mind that our results are an expectation conditioned on that survival. That conditioning is sometimes called survivor bias.

The goal of a real world simulation is to estimate the probability of future outcomes. This usually begins by comparing history to a model—many of

turn to page 22

From Page 21

which are inspired by Markowitz—and decompose historical returns into a risk free return (the return on low duration government bonds, for instance), standard deviation assuming a lognormal evolution of prices, and risk premium. Those parameters are based on a time period that is deemed likely to reflect the simulated period, usually by ignoring historical data before some point in the past.

The major choice is deciding whether the market moves in absolute increments or increments that are relative to yield.

For instance, in 1992 the English pound left the European Monetary Union under pressure and in 1997 the Bank of England achieved operational autonomy. Periods including 1992 probably won't be repeated, while in 1997 consistency of policy was achieved and seemed likely to persist. Calibrations to the U.K. Gilt yield curve start to look similar after 1995, so that may reasonably be the start of the sample calibration period. In the United States, there have been several distinct regimes such as the Volcker Fed period of active inflation fighting, and the Greenspan Fed that has alternated between inflation fighter and liquidity provider. More formal methods, like maximum likelihood estimation, can also identify periods that are not like the recent past, although they do not prove anything about the present.

All these methods rely entirely on the past. But the United States and maybe much of the world is undergoing a shift: an aging population with needs shifting from consumption to saving is combining with the increased productive capacity of globalization in a way that is very likely to change future risk premiums. Capital that was once dear has become relatively cheap, asset prices higher: the risk premium has fallen. Ways of thinking about risk have changed. Most investors now invest against a benchmark that is not cash. In fact, a professional investor's benchmark is a risk free asset, since the risk premium and standard deviation are measured against it. The rise of the benchmarked investor calls into question everything we believe about risk premiums. Prices of U.S. government bonds, in many cases, move inversely with risky assets, thereby acting as "negative risk" assets. That implies that cash indexed investors, such as hedge funds, might be willing to accept a negative risk premium on at least one type of positive standard deviation asset, the U.S. Treasury bond, because of the way it combines with their other holdings. The dynamics of the U.K. Gilt curve also suggest that the highest standard deviation long bond may not generate the highest return over time. Is there reason to believe in a systematic standard deviation linked risk premium when capital is not the constrained asset and cash is not a universal safe haven?

That caveat aside, there are several approaches for taking observed markets and turning them into scenarios that can be used to price financial guarantees. The first is to use the same sort of models frequently used to generate risk neutral scenarios. The major difference is setting the risk premium to some value other than zero, and using the historical volatility of the market rather than deriving a volatility that matches market prices. The major choice is deciding whether the market moves in absolute increments or increments that are relative to yield. Naturally, actual yield curve movements are a combination of both.

Another approach is to use a minimally parametric approach such as that of Rebonato (2003), which simulates yield curve mechanics by sampling the actual yield changes and local curvature from historical data for some period, then jumps stochastically to another time interval. Calibration is limited to picking a time period and jump parameters, and determining the weight each local curvature applies to its associated yield curve point. This model could very likely be extended to other yield curve exposures, such as risky bond spreads or yields.

There are many equity models that reproduce some features of historical equity returns in the United States, such as the volatility of lognormal standard deviations, the correlation of standard deviations with changes in price and the less than linear increase in variance of equity prices with time. There are the regime switching lognormal model and stochastic log volatility model associated with the C-3 Phase 2 effort of the AAA. There is a model described by Ed Thorpe (2004, 2005) using a GARCH approach combined with the Students T distribution under which the 1987 crash is not an impossible outlier. A model no doubt could be constructed along the lines of Rebonato (2005) described above with equity indices substituting for yield curve points.

Most convenient is a model like Heston's (2000) GARCH model for equities that has both a closed form solution for European options, given a historical equity path, and can be used to estimate future risk neutral parameters, such as market implied volatility. That allows for a mixing of the real world and risk neutral pricing by using the real world path to give a best estimate of the risk neutral parameters that allow for the simplified risk neutral calculation. This sounds complex but is simpler than directly defining the evolution of risk neutral distributions over time.

And how can the modeler be sure that extraordinary efforts are rewarded? Option pricing models at least allow for the reproduction of the input market prices as a test for effectiveness. Success at pricing out of sample market observations, improved tracking over time and more stable parameters are a bonus for doing extra design work. Real world models can be tested for statistical similarity to the input set, but this has to be done with the higher order detail in mind.

Full Circle…

Scenarios of all kinds require careful calibration to achieve their intended result. The process starts by creating a clear picture of the important features of the scenarios. A model that supports those features has to be used. Then a sample set of observable prices or data is identified and criteria for what constitutes a good enough match is determined. Imposing a distributional assumption on scenarios may substantially reduce the richness of distributions and produce major differences in the prices of some types of insurance company financial guarantees. This can be reduced by more sophisticated models or by using statistical techniques that don't assume distributions, but even the most sophisticated approaches can't overcome the data limitations of incomplete or contradictory observations. In an incomplete market, the true price of an unobserved security can't be known with certainty, but its sensitivity to the unknown can be systematically reduced once it is understood.

References

Heston, Steve L., and Saikat Nandi. 2000. A Closed Form GARCH Option Valuation Method. *Review of Financial Studies* 13(3): 585-625.

Rebonato, Riccardo, Sukhdeep Mahal, Mark Joshi, and Lars-Dierk Bucholz. 2005. Evolving Yield Curves in the Real-World Measures: A Semi-Parametric Approach. *Journal of Risk* 7(3): 29-62.

Thorp, Ed. 2004. The Distribution of Stock Price Changes—Part One. Wilmott Magazine. Online at *http://wilmott.com/article.cfm?NoCookies=Yes&forumid=1*.

Thorp, Ed. 2005. The Distribution of Stock Price Changes—Part Two. Wilmott Magazine. Online at *http://wilmott.com/article.cfm?NoCookies=Yes&forumid=1*.

David Hopewell, FSA, is senior manager at Ernst & Young LLP in Philadelphia, Penn. He can be reached at david.hopewell@ey.com.