SOCIETY OF ACTUARIES

Article from
Risks & Rewards
August 2019
Issue 75

ISSUE 74 AUGUST 2019

Risks & Rewards

SOCIETY OF ACTUARIES
INVESTMENT SECTION

Beware of Stochastic Model Risk!

By Stephen J. Strommen

Stochastic models have become a staple of actuarial work. In the life insurance and annuity business, they are often used for hedging of long-term guarantees and are needed for financial reporting where "market-consistent" valuations are required and where regulatory reserving mandates their use.

It wasn't always this way. Previous generations of actuaries put values on long-term guarantees by using conservative assumptions. Since there was no large and open market for long-term insurance guarantees, such values were largely a matter of professional judgment.

Around the beginning of the 20th century, Louis Bachelier was the first to apply the mathematics of stochastic processes to the valuation of stock options. His work implied that one might assign a probability to prices in the future based on assumptions made today. Later in the 20th century, Black and Scholes refined these ideas and incorporated the market price of risk to develop the Black-Scholes formula for stock option prices. Since then, similar techniques have been applied to fixed-income instruments and interest rate derivatives. These techniques have been widely adopted by actuaries and others for valuation of all sorts of out-of-the money options and guarantees. These techniques improve upon previous methods that were less quantitative and based largely on judgment.

Stochastic techniques have been successful at least partly due to their ability to explain market prices. One can choose an applicable stochastic model and fit the

CONTINUED ON PAGE 4

parameters of the model to some market prices. Then those same fitted parameters can be used in the stochastic model to determine a market-consistent price of something that is not actively traded, such as a life insurance or annuity contract with long-term guarantees. This idea has been extended to imply that such models can be used to determine the probability of failure of an insurance company or block of business. The 99.9 percent probability threshold in Solvency II is based on this extension of the technique, as are the probability levels specified by the NAIC for certain reserve and capital requirements.

Unfortunately, blind reliance on such models can be disastrous. Recall the fate of Long-Term Capital Management (LTCM), a hedge fund management firm that applied these models in a big way. The firm's strategy was basically to use market-consistent valuation to identify securities whose actual market price deviated from the market-consistent price, on the theory that the price would converge to become more market-consistent over time. The firm's results in the first few years were stellar. Then in 1998, the firm lost $4.6 billion in a few months and required a $3.6 billion bailout funded by 16 big banks under the supervision of the Federal Reserve. Actual market prices did not behave in the manner their models anticipated, and financial disaster ensued.

I am concerned that many actuaries do not understand the degree of model risk that is present any time stochastic models are used. Just like the founders of LTCM, some actuaries give undue deference to results from a stochastic model. Different stochastic models of the same business produce different results, and the differences can have big consequences.

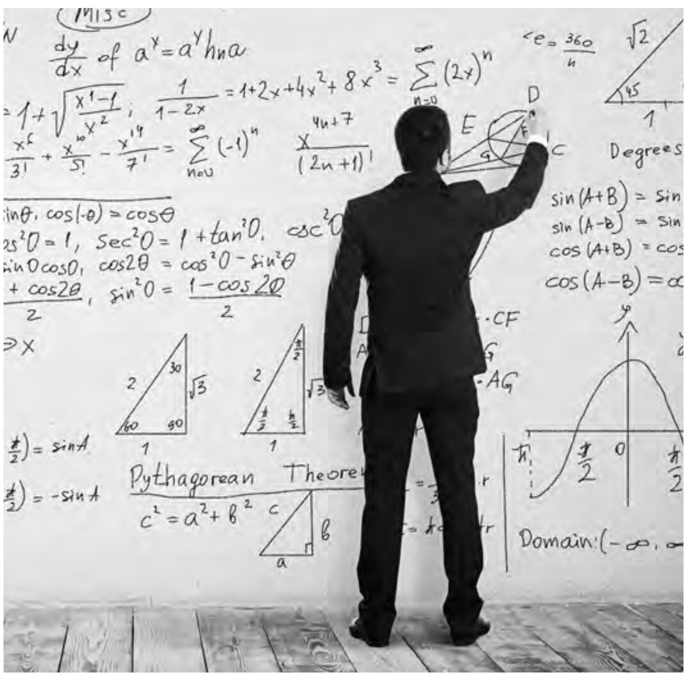

With that in mind, the remainder of this article highlights several areas where model risk arises due to the choice of a stochastic model and its calibration. We focus here on models for future interest rates and equity market returns. Consider the model risk arising from each of the following.

USE OF THE NORMAL DISTRIBUTION

The familiar bell-shaped curve of the Normal or Gaussian distribution is widely used as a mathematical model for uncertainty. It is mathematically straightforward and facilitates derivation of closed-form formulas for many commonly used measures of risk and value.

Unfortunately, the Normal distribution is just a first-order approximate model for risks in the real world. It is well documented that the actual variability of most economic variables is better characterized by a distribution with both fatter tails and a stronger central peak than the Normal bell-shaped curve.

The Black-Scholes formula for stock option prices is one of the most common tools built on the Normal distribution. The "volatility" parameter of the formula is analogous to the standard deviation of the Normal distribution. If the model fit well, then a single value for volatility would approximately fit all market prices. But it does not. The fitting of actual market prices using the Normal distribution results in an "implied volatility surface," which is an array of different values depending on strike price and tenor. The knowledgeable actuary will understand this as evidence that the underlying model does not fit very well. In particular, it is not ideal for use in generating future scenarios for stochastic simulations because a generator can use only a single value for volatility at a point in time, not an array of fitted values.

There are several ways to address this issue when choosing a stochastic model for use in a scenario generator. The three most common are:

- **Stochastic volatility.** The Normal distribution is still used, but the volatility parameter is made to follow its own mean-reverting stochastic process over time. When the volatility is lower than average in the scenario, values clump toward the center of the distribution. When the volatility is higher than average, relatively more tail values are generated. Overall, the ultimate distribution has longer tails and a stronger central peak.

- **Regime switching.** The Normal distribution is still used but the model switches between two regimes, which are characterized by different sets of parameter values for both the volatility and the mean. There is a high-volatility regime and a low-volatility regime, typically with different mean values. Switching between regimes results in an ultimate blended distribution that can have longer tails and a stronger central peak.

- **Different underlying distribution.** The Normal distribution is abandoned as a model of variability within each time step. A different distribution that has longer tails is used instead. There are many choices for such a distribution.

DEALING WITH THE ZERO LOWER BOUND

Interest rates may generally follow a random walk, but the ability to simply hoard cash makes it economically difficult for interest rates to fall much below zero. Slightly below zero is possible due to the expense and risk associated with hoarding cash, but far below zero is arguably not possible while markets continue to function. One would think that any interest rate model in common use would need to reflect this near-zero lower bound on interest rates.

Not so. For example, the Ho-Lee lattice model is commonly used for valuation of callable bonds and other fixed-income instruments with options. The underlying model is a recombining lattice for paths of future interest rates, with equally spaced up and down jumps. When carried far enough into the future, some paths through such a lattice involve negative interest rates. Yet this model is in common use because of its speed and efficiency and mathematical tractability.

Several approaches for dealing with the zero lower bound are in circulation. Among them are these:

- **Make the volatility of interest rates proportional to the current interest rate.** When interest rates are low, volatility becomes low so that it becomes unlikely that a random shock will push interest rates below zero. When interest rates are high, volatility is high, as happened in the early 1980s in the U.S. This approach has an effect on the implied future distribution of interest rates, making it skewed with a longer tail on the high side. Two versions of this are in common use.

  - The Cox-Ingersoll-Ross method makes the volatility proportional to the square root of the interest rate. The ultimate distribution of future interest rates tends toward a noncentral Chi-Square distribution.

  - The Black-Karasinsky method uses a constant volatility for the log of the interest rate rather than for the interest rate itself. The ultimate distribution of future interest rates tends toward a lognormal distribution.

- **Impose a zero floor at each time step.** If the stochastic process produces an interest rate below zero at any time step, set it to zero before proceeding to the next time step.

- **Track the theoretical path of the interest rate separately from the lower bound.** Under this method, the stochastic process is allowed to take its course and produce negative interest rates, but the interest rates actually output from the generator are floored at zero. This can lead to scenarios with extended periods of very low interest rates.

PERSISTENCY

For the sake of discussion, let's accept the proposition that interest rates are mean-reverting. In the U.S. the Federal Reserve largely controls interest rates. The Fed has a target level and moves interest rates up or down relative to that target depending on whether economic stimulus or inflation control is more important at the moment. A mean-reverting random walk seems to be a reasonable stochastic model for interest rates in these circumstances.

> Different stochastic models of the same business produce different results, and the differences can have big consequences.

In recent years, there has been concern that a simple mean-reverting random walk may not be particularly realistic. While interest rates may revert to the mean in the long run, they have tended to be very persistent and stay within a narrow range in the short run. This suggests a stochastic process with persistence, whereby scenarios can remain far from the ultimate mean for long periods of time. A simple mean-reverting model is anti-persistent because any diversion from the ultimate mean is immediately countered with a stochastic tendency to revert back to the ultimate mean.

Persistence can be increased by modifying the stochastic process. In a simple mean-reverting model, the mean is constant. In more complex models, the mean itself can be made to vary over time.

- In a double-mean-reverting model, there is a current mean and a long-run mean. The long-run mean is constant, but the current mean follows a simple mean-reverting process. In the short run, scenarios revert to the current mean, not to the long-run mean.

- In a regime-switching model, there are two different values for the mean. Only one is active for each time step. There is a probability of switching from one to the other at each time step, but that probability is typically low.

I have measured the proportion of 30-year scenarios containing periods where short-term interest rates remain below 2 percent for 10 years or more. Similarly calibrated generators (based on the distribution of interest rates 30 years in the future) using different stochastic processes yielded proportions that varied from less than 3 percent to more than 10 percent. Such differences could easily affect modeled capital requirements in connection with long-term minimum interest guarantees.

CHOICE OF CALIBRATION PERIOD

Calibration of an economic scenario generator for stochastic simulations is typically done using historical data over an extended period of time. Unfortunately, the historical record does not include enough time steps to provide stable calibration. The choice of time period to use for calibration can affect the results significantly, making calibration unstable. It can be instructive to compare the results of stochastic simulations using alternate calibration periods for the parameters of the scenario generator.

This should be kept in mind any time market-consistent or "risk-neutral" scenarios are used. Such scenarios are typically calibrated to market conditions on a single day. When used for stochastic simulations that extend over decades, the results can be notoriously unstable unless there is some sort of mean reversion built into the parameters of the stochastic process over time.

CALIBRATION FOR REGULATORY PURPOSES

Sometimes the parameters of a generator may need adjustment in order to meet regulatory calibration targets. This should be viewed as substituting the regulator's judgment for one's own. The regulatory calibration targets may be based on a different calibration period, a different stochastic model or both. When employing stochastic modeling, one should not blindly accept the regulator's judgment as embedded in regulatory calibration criteria. Valuable insight can be gained by running a stochastic model using your own generator and your own calibration before determining the effect of any adjustment needed to meet regulatory calibration requirements.

CONNECTIONS BETWEEN INTEREST RATES AND EQUITY RETURNS

When stochastic economic scenarios were first proposed for actuarial risk management, the theoretical work on scenario generators accepted certain long-term economic relationships as axiomatic. Interest rates and inflation were related. Stock returns and inflation were related. Stock returns were volatile but had an expected mean higher than interest rates due to a risk premium. The Wilkie model, an early stochastic scenario model, reflected these axiomatic relationships.

Recently, a more statistical approach has taken hold, and relationships that cannot be proven as statistically significant based on historical data are often abandoned. In particular, the relation between interest rates and equity returns is often treated as nonexistent because it cannot be proven beyond statistical doubt. The idea that the limited time span of the modern historical record provides insufficient data to either prove or disprove such conjecture tends not to be considered. An example of this is the generator currently mandated by the NAIC for regulatory use, which treats interest rates and equity returns as independent.

If you believe in a relationship between risk and expected return, then you may accept it as an axiom that in a real-world model, the expected return on equities should exceed the risk-free rate by an expected risk premium at every time step in every scenario. In a risk-neutral model, the expected return on equities should equal the risk-free rate and the risk premium should be zero. Such relationships between risk and return are violated when the stochastic processes for interest rates and equity returns are independent.

Just because something cannot be proven does not mean it isn't true. This is true in statistics, it's been proven true by Gödel in mathematics, and I believe it is true in the context of stochastic economic scenario models. As a matter of professional judgment, one should be careful when using models that do not reflect relationships that one believes to be true.

SUMMARY

The point of this article is not to discredit stochastic models. Such models can be very useful tools for analysis of risk. The point here is that results from a stochastic model should not be given any more deference or be considered more exact than any other kind of actuarial estimate. Instead, they should be viewed as approximate guidance that can best be used to inform professional judgment. The choice of model and its calibration should be treated with just as much care and review as the underlying actuarial assumptions in a deterministic calculation.

As was demonstrated by the case of LTCM, blind reliance on stochastic models can pose a significant risk. Careful review and appropriate use of such models can lead to rewards. So this fits squarely within the risk manager's purview of risks and rewards! ■

Stephen J. Strommen, FSA, CERA, MAAA, is an independent consultant and owner of Blufftop LLC. He can be contacted at *stevestrommen@blufftop.com*.