SOCIETY OF ACTUARIES

Article from
The Financial Reporter
June 2019
Issue 117

Common Misunderstandings of Risk-neutral Valuation

By Stephen J. Strommen

One prominent idea in the recent development of accounting for insurance contracts is the immediate recognition of profit or loss due to changes in market values of assets and liabilities. Implementation of this idea requires determination of market value not just for invested assets but also for insurance liabilities. While most invested assets are traded in a market, most insurance contracts are not, so there is no easy way to obtain a "market value" for most insurance contracts. The idea of market-consistent valuation has gained traction to satisfy this need, and stochastic risk-neutral valuation has come to the fore as a widely recognized approach to market-consistent valuation.

As an actuary involved in discussions of new accounting standards, I have encountered several misconceptions about risk-neutral valuation, even among some experienced and prominent financial reporting actuaries and regulators. This article highlights several of these misunderstandings with an eye toward putting the debate in this area on a more scientific basis.

There is a common understanding of the following basics of stochastic risk-neutral valuation:

1. The time value of money is characterized by the short-term (single period) default-free interest rate (the "short-term risk-free rate").

2. The future path of the short-term rate is uncertain and can be characterized by a random walk or other stochastic process.

3. In risk-neutral stochastic valuation, the random walk or stochastic process governing the future path of the short-term rate is calibrated so that:

    a. the expected or central path of the short-term rate is the forward rate path of the observed risk-free yield curve; and

    b. the volatility is such that market prices of options and other derivatives are reproduced.

The mathematical justification of risk-neutral stochastic valuation is complex. While many actuaries understand the three points just mentioned, I have often encountered the following misunderstandings regarding their implications.

MISUNDERSTANDING 1: THE MARKET'S EXPECTATION OF FUTURE SHORT-TERM RATES IS EQUAL TO THE FORWARD RATE PATH OF THE OBSERVED RISK-FREE YIELD CURVE

Point 3.a. above says that risk-neutral scenarios are calibrated so that the expected path of the short-term rate equals the forward rate path of the risk-free yield curve. So it is true that the risk-neutral expectation of future short-term rates is equal to the forward rate path of the observed risk-free yield curve. But the market's expectation is not the same as the risk-neutral expectation. The probability distributions of future events and their expected values differ between the real world and the risk-neutral world. The real-world distributions are referred to as the P measure, and the risk-neutral distributions are referred to as the Q measure. The expected path under the P measure is different from that under the Q measure.

MISUNDERSTANDING 2: THE EXPECTED PATH OF THE SHORT-TERM RISK-FREE RATE IS HIGHER IN REAL-WORLD SCENARIOS THAN IN RISK-NEUTRAL SCENARIOS

Actually, the expected future path of the short-term risk-free rate is **lower** in properly calibrated real-world scenarios than in risk-neutral scenarios.

This misunderstanding probably arises because of the way equity investments are simulated when risk-neutral scenarios are used for simulation. In a risk-neutral simulation, the distribution of equity returns is centered on the short-term risk-free rate—even though the expectation in the real world is that, on average, equities will earn a higher return that includes a risk premium. Basically, the projected cash flows from equity investments are lower in a risk-neutral simulation than in a real-world simulation.

Fixed-income securities are treated differently than equities. The cash flows are fixed, so it is the discounting of those cash flows that must be different.

To understand why the discounting is different, one must understand the nature of the "risk-free" yield curve. Only the short-term rate is risk-free. All longer-term rates involve lock-in of an interest rate in an environment where interest rates can change. Lock-in is a risk to the investor because interest rates could rise, resulting in a loss of market value. That risk has a price, and it is included in the risk-free yield curve in the form of term premiums. Long-term risk-free rates are normally higher than short-term rates because of the existence of term premiums, which are a form of risk premium.

Since the risk-free yield curve includes term premiums, the forward rates in that curve include term premiums. To get the market's expectation of the path of the short-term rate, those term premiums must be removed. The market's expectation of the future path of the short-term rate is lower than the path of forward rates in the risk-free yield curve because the term premiums are removed from the long-term forward rates to get the expectation for the short-term rate.

Term premiums are not insignificant. For example, the real-world stochastic interest rate generator mandated by the NAIC for use in VM-20 valuations has a parameter to set the average term premium in the 20-year rate 100 basis points higher than that in the one-year rate.

Term premiums increase by length of time from the valuation date. The longer the scenarios, the greater the difference between risk-neutral and real-world scenario paths. This should be an important consideration when using risk-neutral valuation for long-term insurance contracts. In the investment world, the risk-neutral approach is primarily used for valuation of comparatively short-term derivative securities where the difference between real-world and risk-neutral scenario paths is much smaller. Extension of the risk-neutral approach to much longer-term contracts is somewhat akin to extending the results of a linear regression to points far outside the sample used to calibrate the regression. This is especially true when extending risk-neutral valuation to contracts that last beyond the end of the observable yield curve.

MISUNDERSTANDING 3: ONLY A RISK-NEUTRAL VALUATION CAN BE MARKET-CONSISTENT, SO REAL-WORLD SCENARIOS SHOULD NOT BE USED FOR MARKET-CONSISTENT VALUATION

There is a common misunderstanding that the terms "market-consistent" and "risk-neutral" mean the same thing in the context of valuation. In fact, risk-neutral valuation is just one approach to performing a market-consistent valuation.

This misunderstanding may have arisen partly because many "real-world" scenario generators are not market-consistent. In order to be market-consistent, a generator must be calibrated to current market conditions on the scenario starting date. Many real-world generators are used to measure capital adequacy and are not frequently recalibrated because they are not used for valuation. The focus for their use is the outlier scenarios, not the central scenarios that get most weight in a valuation, so calibration of the central scenarios is not important.

Nevertheless, a real-world scenario generator can be market-consistent if it is calibrated on each valuation date. Three aspects of current market conditions must be included in the calibration:

a. The expected path of future short-term interest rates, based on the yield curve with term premiums removed

b. The volatility of interest rates, based on the market prices of derivatives

c. The market price of risk

The market price of risk is not directly observable, and neither are the term premiums. They can be inferred indirectly using a combination of historical data and current prices. Risk-neutral calibration gets around this problem by treating the market price of risk and term premiums as zero and adjusting the expected path and volatility to compensate. The theory that justifies that is complex, but the basic idea is that the market price of risk becomes implicit in the adjusted path and volatility of future interest rates in risk-neutral scenarios.

Real-world calibration is sometimes criticized because it requires explicit treatment of the market price of risk and is, therefore,

There is a common misunderstanding that the terms "market-consistent" and "risk-neutral" mean the same thing.

more subject to judgment. This is based on the misconception that risk-neutral valuation does not involve judgment, which will be addressed next.

MISUNDERSTANDING 4: CALIBRATION OF RISK-NEUTRAL SCENARIOS IS OBSERVATION-DRIVEN AND INVOLVES LITTLE JUDGMENT

Risk-neutral scenario calibration is rooted firmly in observed data. But significant judgments are still involved.

The first judgment is the choice of underlying stochastic process to be calibrated. For interest rates, there are one-factor models, two-factor models, stochastic volatility models, regime-switching models, zero lower bound models, and so on. The stochastic shocks in these models can be normal or lognormal or can use other distributions. The choice of stochastic process will affect characteristics of the generated scenario set, such as the frequency and length of periods of persistent low interest rates. These characteristics can certainly affect the valuation of insurance contracts, especially those with minimum interest crediting guarantees.

The second judgment is the choice of volatility to use when generating stochastic scenarios. Calibration will provide a volatility surface—that is, a range of implied volatilities that vary by strike price and tenor. This range of implied volatilities is an indication that the model does not fit perfectly, but that point is often passed over. When generating stochastic scenarios, volatility can have only one value in each time step, not a different value for each strike price and tenor, so judgment is necessary in selecting the volatility to use.

The third judgment is the measurement of the risk-free rate. For valuation of insurance contracts that are illiquid, it is generally accepted that the observed yield curve for U.S. Treasurys is inappropriate because Treasurys are very liquid. Illiquid securities have higher yields than liquid securities, so an "illiquid default-free" yield curve is suggested for use. Such a yield curve can be U.S. Treasurys plus an illiquidity adjustment. Sometimes the illiquidity adjustment is given other names, such as a matching adjustment. Whatever the name, setting the size of the adjustment requires judgment, and there is significant debate over the appropriate size of adjustment to be made when valuing different kinds of insurance contracts.

The three judgments listed here can significantly affect the results of risk-neutral valuation for long-term insurance contracts. In my view, these judgments within the risk-neutral approach are just as significant as the judgments required in the real-world approach to market-consistent valuation.

MISUNDERSTANDING 5: THE MARKET PRICE OF RISK IS THE SAME FOR EVERYONE

One important aspect of the theory behind risk-neutral valuation is that the market price of risk is a single figure and is the same for everyone. Calibration of risk-neutral scenarios does not quantify the market price of risk but builds it in implicitly through the expected path and volatility of future interest rates.

Real-world market-consistent valuation requires one to specify the market price of risk. Often that is done by equating the market price of risk to the cost of capital. In the real world, we know that the cost of capital is not the same for everyone.

The fact that the market price of risk is not the same for everyone is fundamental to the very existence of the insurance business. Understanding this provides some insight into the debate over determination of the appropriate discount rate for market-consistent valuation of insurance contracts.

The difference between parties for the price of risk can be considerable. Let's define the price of risk as the cost of keeping available the amount of money needed to be made whole after a risk event occurs—that is, keeping money available to pay for the potential loss. Consider a family that owns its home. It must bear the risk of destruction of the home through fire or other disaster. In the absence of risk sharing, the amount they must keep available to restore their home in the event of loss is the full value of the home plus the cost of potential temporary housing. In the absence of risk sharing, that is the price of bearing the risk.

With insurance, the cost of bearing that risk can be vastly reduced to the size of a small annual homeowners insurance premium because that is all that's required to make available the money required to replace the family home if it is destroyed. The cost of the risk to the insurer is much lower than to the family because the insurer makes use of risk sharing.

Basically, the financial purpose of insurance companies is to reduce the market price of insurance risk through risk pooling and diversification. To accomplish this, insurers are motivated to increase in size (to increase risk pooling) and to diversify (to reduce correlation of risks). As a result of these activities, the price of risk for insurers is reduced. Insurers can provide risk protection with what amounts to a lower cost of production and can, therefore, sell it at a low price.

The fact that the market price of risk is not the same for everyone is fundamental to the very existence of the insurance business.

This applies not only to insurance risks but also to investment risks, such as bond defaults. The price of this risk is reduced for insurers precisely because of pooling and diversification with other risks. This means that the expected net investment return for the insurer, after subtracting the insurer's price of risk, is higher than the so-called risk-free rate.

I understand that the prior paragraph is heresy to some economists and actuaries. But when you think about it, the concept at work here is the same as that which suggests that introduction of technology that lowers the cost of production for a manufactured good will lead to lower market prices.

To continue with this heresy, consider the idea that insurers pass their investment returns on to customers through the pricing of insurance products. A simple example is the pricing of lifetime income annuities. Insurers typically back annuities by investments in a portfolio of defaultable bonds. Their low cost for bearing the default risk is passed on in the competitive marketplace by pricing with net investment returns higher than the risk-free rate. That's because their expected investment return—net of defaults, expenses, and net of the cost of capital—is significantly greater than the risk-free rate. (Challenge to the reader: Do the math. See sidebar, pg. 18.) Call the excess over the risk-free rate a liquidity adjustment or a matching adjustment or something else, but I believe it comes partly from pooling and diversification, not just liquidity.

I believe the liquidity adjustment or matching adjustment is required for a risk-neutral valuation to be market-consistent. This is based on observation of real market prices. Those who push back on this sometimes argue that life income annuities are often mispriced by insurance companies; the market prices are too low because the investment return assumptions exceed the risk-free rate. I find that argument to violate the scientific method. In science, observations take precedence over predictions based on theory. Those who say annuities are mispriced because of such investment return assumptions give predictions of their theory precedence over observations of actual market prices.

CONCLUSION

The risk-neutral approach to valuation has come to the fore in recent years as accounting standards have moved toward use of market-consistent valuation. Actuarial standards are now being drafted regarding compliance with the new accounting standards. In drafting these standards, some have suggested that the risk-neutral approach should be required for market-consistent valuation. This article has highlighted some misunderstandings about the risk-neutral approach that have come up in such discussions, with the hope that better understanding will lead to standards that reflect the complexity of the issue and allow alternate methods and professional judgment where appropriate. ■

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

DO THE MATH

What is the market-consistent discount rate for valuation of a lifetime payout annuity by an insurance company? How does it compare with the risk-free rate?

Assume that the insurer invests in A-grade corporate bonds. For simplicity, we look at the net spread on a 10-year A-grade corporate bond.

Gross credit spread:	133 bps (source: NAIC tables for VM-20 valuation)
Less:	
Expected defaults	18 bps (source: NAIC tables for VM-20 valuation)
Investment expenses	10 bps
Cost of capital	48 bps (8% capital requirement x 6% cost of capital rate)
Net spread:	57 bps

Based on these assumptions, the market-consistent valuation uses a discount rate that includes a 57 basis point spread over the risk-free rate. This is a bit oversimplified because the calculation should reflect a weighting of net spreads at different points on the yield curve, assuming the insurer would purchase an array of bonds to match the expected cash flows of the annuity. And the cost of capital is an estimate that involves judgment. Nevertheless, the market-consistent spread over the risk-free rate is significant, because a reasonable estimate of the insurer's cost of capital is much less than the market credit spread.