Commentary on the Conning interest rate model

Autumn 2022

Background and intended purpose

Scenario generators have come into wide use for risk management and valuation of insurance contracts. Modern valuation techniques both in the US and internationally often employ stochastic scenarios to calculate the financial statement value and capital required for insurance contracts. The practical purpose is to make the valuation of insurance liabilities more consistent with the valuation of the assets that back them. Consistent valuation avoids a valuation mismatch and provides a better measure of financial condition.

While this common idea has been employed both in the US and internationally, there is an important difference in the way it has been implemented. The difference arises because assets are valued differently under different accounting standards. Under the international standard (IFRS 17) assets are held at market value. Under US statutory accounting, most assets are held at amortized cost or book value. When assets are valued differently, a liability valuation consistent with asset valuation must be done differently.

Scenario generators have been developed to support these different valuations. But, just as the valuations are different, the characteristics of an optimal generator differ depending on the kind of valuation. A generator that is well-suited for one purpose may be ill-suited to another. That is the central issue with use of the Conning interest rate generator for US statutory valuation. That generator is well-suited to market-consistent valuation as under IFRS 17 or Solvency II. But it is ill-suited to US statutory valuation. The main difference is that valuations under IFRS 17 and Solvency II are intended to be consistent with market values and typically use risk-neutral scenarios, while US statutory valuation is intended to be consistent with asset book values and requires use of real-world scenarios. The same characteristics that make the Conning generator good for risk-neutral scenarios make it ill-suited to real-world valuation.

Before getting into technical details, consider the difference in basic purpose for risk-neutral and real-world scenarios. Risk neutral techniques were developed for the purpose of market-consistent valuation on a single date. The goal for their development is use of observable information about current market prices for assets. It is not necessary for the scenarios to simulate realistic movement over time. In fact, the movement of future prices using risk-neutral calibration is known not to be realistic. Real-world techniques were developed for simulation over time, generally long periods of time. There the evolution of prices over time is the main focus. This difference in purpose leads to development of different mathematical forms of generators. A mathematical form that is well-suited to the risk-neutral context may be poorly suited to the real-world context. That is the case with the Conning interest rate generator.

Another consideration has arisen in recent years, after the 2008 financial crisis. Before then it was generally accepted that interest rates fluctuated in a range above zero. However, in recent years central banks have pushed rates to zero or even below in some cases. The behavior of interest rates around zero is different from behavior well above zero, and generators developed in the past are not equipped to deal with this new behavior. Conning's generator does not handle behavior at low interest rates well.

Technical details

To understand these aspects of Conning's generator one must evaluate it at a very technical level. One must understand exactly what the formula is, how it is used, and how it behaves in the context of scenario generation, particularly when interest rates are low. We begin by reviewing the formula and the way it is used in a risk-neutral context. It will be seen to have advantages there. Then we review how it is used in a real-world context and explain how some features that represent strengths in the risk-neutral context create weaknesses in real-world application.

Conning's model is a 3-factor Cox-Ingersoll-Ross model where the short-term interest rate is the sum of three state variables (each of which follows a one-factor Cox-Ingersoll-Ross model) and a constant referred to as the "shift". It is what is termed a "short-rate model" because the remainder of the yield curve (other than the short-term rate) is defined as an expectation based on the future path of the short-term rate. More formally:

Let r_t be the short-term rate at time t. In Conning's model, the stochastic process for this short-term rate is:

$$r_t = f_{1,t} + f_{2,t} + f_{3,t} - s$$

where f_1, f_2, f_3 are three independent one-factor Cox-Ingersoll-Ross processes, each with their own mean reversion point, mean reversion speed, and volatility, and s is a constant subtracted from their sum.

In a short-rate model the yield curve beyond the short-term rate is completed by calculating spot prices and the corresponding spot rates using the following formula for spot prices:

$$P_T = E^Q\left[exp\left(-\int_0^T r_t dt\right)\right]$$

The Cox-Ingersoll-Ross process used for each of the three factors cannot go below zero. The shift s is therefore what allows this process to produce negative interest rates.

Advantage 1: The model can produce negative interest rates.

In this model the three factors f_1, f_2, f_3 tend to behave like the three components of the yield curve as derived from principal component analysis. The three components can be described as level, slope, and curvature. Together, three components like this can be used to produce a great variety of yield curve shapes. Calibration of the parameters for use in risk-neutral valuation tends to maximize the ability to produce the variety of yield curve shapes that have been observed.

Advantage 2: The model can produce a wide range of yield curve shapes.

As mentioned above, the completion of the yield curve is done using a statistical expectation using the stochastic path of the short-term rate. That process for completing the yield curve makes this an arbitrage-free model.

Advantage 3: The model is arbitrage-free.

Not all models allow solving for a closed-form expression for the statistical expectation needed to complete the yield curve in this way. When such an expression is not available, numerical simulation is needed to calculate the expectation, and that is time-consuming. This model does have such a closed-form expression.

Advantage 4: The model provides a closed-form expression for completing the yield curve.

Clearly, this model has advantages for calibration and use in a risk-neutral context. Risk-neutral calibration focuses on setting the parameters in a way that allows the model to produce the widest variety of yield curve shapes. That's done because the shape of the yield curve defines market prices at a point in time, and the primary goal of risk-neutral calibration is consistency with market prices at a point in time – the valuation date.

Risk-neutral valuation is designed to produce a market-consistent value of something purchased in the past which provides the owner with cash flows in the future – cash flows that may be uncertain in timing or amount and which can be modeled stochastically. That includes a large portion of insurance contracts. Valuation is done by projecting the contract cash flows along each scenario in a risk-neutral set, using path-wise discounting to get the present value in each scenario and then averaging the results.

This works well when contract cash flows can be projected without reference to future returns on related invested assets. But when the insurance contract cash flows depend on future returns on invested assets, there is a problem. Risk-neutral scenarios do not provide realistic future returns – they are designed and calibrated to produce market-consistent values only on the scenario start date. As for the future, they employ "risk-neutral probabilities" which are known to be different from "real-world probabilities". In technical terms, risk-neutral scenarios use the Q measure while real-world scenarios use the P measure, and different measures are different sets of probabilities for future returns.

To solve this problem, one can use real-world scenarios rather than risk-neutral scenarios. Real-world scenarios provide realistic paths of future investment returns. Real-world scenarios also make sense when assets are held at something other than market value. In that situation, for the valuation of liabilities to be consistent with the valuation of assets, the discount rate used must be based on the expected future return on the assets as measured under the accounting standard in use. That's why the principle-based valuation approach in the US uses real-world scenarios, not risk-neutral ones, and uses the investment return on supporting assets (as measured using statutory accounting) as the discount rate for future cash flows.

When using Conning's interest rate model in a real-world context, we encounter some problems. They occur mainly due to differences in the goals of calibration under the real-world approach, but also due to behavior of this model when interest rates approach zero.

Problem 1: Historical movement of state variables exhibits strong correlation, but the model requires independence. The model is incapable of simulating realistic correlated movement of the state variables, and that means it cannot simulate realistic paths of yield curves through time.

Real-world calibration of an interest rate model focuses on realistic movement of the state variables over time. Given any calibration of this model, one can perform a historical exercise of fitting the state variables to historical yield curves. In real-world calibration one goal is to have the historical movement of the state variables be similar to that which could be produced by the model. When this is done using Conning's model, one result is that historical movement of the state variables is significantly correlated. But the model makes the state variables move independently. Independence is required to obtain a closed-form formula for completing the yield curve. The closed-form formula is therefore an advantage in the risk-neutral context which creates a disadvantage in the real-world context where the path through time matters.

Problem 2: The model cannot provide realistic mean reversion at the same time it provides good fit to yield curve shapes. This is a conflict between real-world and risk-neutral calibration.

On aspect of real-world movement is the rate of mean reversion. Mean reversion of interest rates is slow but must be present in a model to prevent interest rates from rising or falling indefinitely; there is a normal range within which interest rates fluctuate. The problem here is that calibration of Conning's model focusing on fitting yield curve shapes results in mean reversion rate near zero in one factor. Raising that mean reversion rate to a real-world level significantly reduces the ability of the model to reproduce historical yield curve shapes.

The model has two sets of parameters, one labeled risk-neutral and one labeled real-world. They are understood to be different, and that difference has been used in Conning's calibration to solve problem 2. The real-world mean reversion rate for one factor is significantly higher than the risk-neutral mean reversion rate. But that creates another problem.

Problem 3: Conning's calibration embeds unrealistic term premiums.

One advantage of the model is that it is arbitrage-free. In an arbitrage-free model with both real-world and risk-neutral calibrations, the difference between them is an implicit market price of risk.

The market price of risk is an important element of calibration of a real-world model. When calibrating a model for equity returns, the Sharpe ratio is often a key measure. The Sharpe ratio is the ratio of expected risk premium to volatility. That is exactly the same concept as the market price of risk, and it has a place in evaluating interest rate model calibration just as much as in evaluating an equity model calibration.

The problem is that the market price of risk implied by the difference between Conning's risk-neutral and real-world parameters is unreasonable, and that arises largely due to the significant difference in mean reversion rate mentioned above. Embedding an unreasonable assumption for the market price of risk raises a significant objection to any claim that the model represents real-world market behavior.

The implicit market price of risk can be illustrated by calculating the implied term premiums in the yield curve. Term premiums are the reward investors demand for the risk of locking in a fixed rate for a period of time. Fixed income securities have price volatility due to fluctuating interest rates. Term premiums are the reward for that volatility and are in theory proportional to it. Since price volatility is greater for longer term fixed income securities, term premiums are higher for longer maturities, where a typical level is between 1% and 2%.

The implied term premiums can be calculated directly. Starting from the values of the state variables at a point in time, complete the yield curve once using the risk-neutral parameters and once using the real-world parameters. Subtract the real-world curve from the risk-neutral curve. The difference is a set of term premiums by duration. One would expect them to start at zero for money market funds and increase for longer maturities to between 1% and 2%, being consistently positive.

Here is a chart of the term premiums for some initial yield curves, using Conning's calibration:

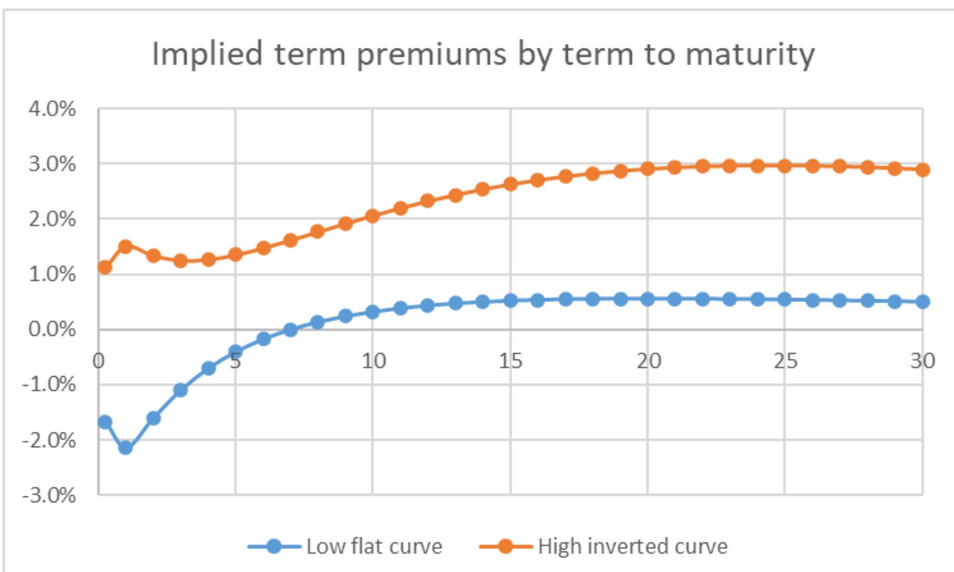

When fit to a low flat curve like those in 2020, the term premiums on the early part of the curve are negative, implying a negative market price of risk. That's analogous to assuming equity investments earn less than the risk-free rate on average.

When fit to a high inverted curve like those in the 1980's, the term premiums are unreasonably high on the long end of the curve.

While not shown here, the term premiums on the short end of the curve are extremely sensitive to the shape of the curve to which the state variables are initialized.

All these issues arise from using substantially different mean reversion rates in real-world vs. risk-neutral parameters.

Problems 2 and 3 are closely connected. To solve problem 2 one needs significantly different mean reversion rates in the real-world and risk-neutral parameter sets. But that is exactly what causes problem 3. If you solve problem 2 you get problem 3. If you solve problem 3 you get problem 2. There is no way to calibrate the model to get around this dilemma.

The Academy of Actuaries provided an alternate calibration representing a compromise. That just means each problem was reduced but neither was eliminated. Even those who provided the compromise are not satisfied with it.

Problem 4: The model requires use of a floor, and each kind of floor introduces unfavorable baggage.

Interest rates near zero create two more problems. The first revolves around how low interest rates can go. Recall that one parameter of the model is a “shift” to allow negative interest rates. The lowest negative interest rate the model can produce is negative by the amount of the shift. This is a problem because calibration of this model to historical data, whether risk-neutral or real-world in nature, results in a very large shift on the order of 10% or more. The result is a generator that produces interest rates that go too far negative too often. This has been observed in all scenarios produced by the model and has been addressed by modifying the scenarios after they have been generated, using a floor.

Adding a floor changes the model. Each of the two kinds of floors introduce issues. The “Generalized Fractional Floor” results in scenarios that are no longer arbitrage-free. The “Shadow Rate Floor” retains the arbitrage-free property in an approximate way but introduces a time-consuming process to create a conversion table whenever the parameters of the model are changed.

Problem 5: The model poorly fits recently observed low yield curves.

The other problem with interest rates near zero is that while the model can technically fit low historical yield curves, that fit is subjectively very poor. Fitting the model to a yield curve is said to be possible if the model can reproduce three points on the curve exactly. But a fit can be possible while still being poor. Here is the model’s fit to the 2020 year-end yield curve:

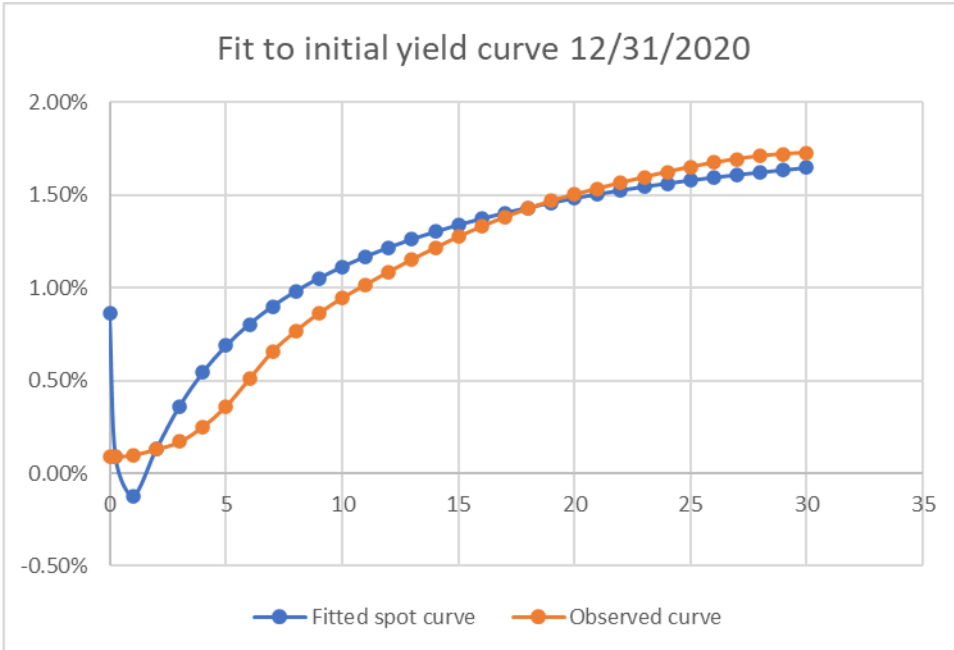

The odd shape of the fitted curve indicates the model does not well reproduce the kind of yield curve shape observed when interest rates are low. In this case, the poor fit at the short end of the yield curve is what causes the unrealistic behavior of the short end of the curve that has repeatedly been observed in generated scenarios.

Problem 6: There is no stochastic volatility, thereby underrepresenting that aspect of interest rate risk.

One stylized fact about interest rate behavior is that volatility fluctuates. Interest rate risk is greater when volatility is higher, and that occurs from time to time. This can be simulated using stochastic volatility. In the Conning interest rate model volatility is constant, so there is no reflection of this stylized fact in the model.

Two other problems arise because of NAIC directives for use of the model and Conning's approach to implementing them.

Problem 7: NAIC requirements regarding the mean reversion point and its criterion for satisfactory "low-for-long" behavior put extreme stress on the calibration of the model, thereby increasing the severity of other issues.

The NAIC developed its formula for the mean reversion point in the context of a different model than Conning's. Different models produce different distributions around the mean reversion point, and it is the tails of the distribution rather than the mean reversion point that are used to calculate reserves and capital. Because of this, the mean reversion point in use for the old generator should be revisited for use with Conning's generator or any other generator. In the alternate calibration provided by the AAA a higher mean reversion point was used for this reason.

The low-for-long criterion set by the NAIC is that 5% of scenarios should have average rates over the first 30 years that are less than the starting rate, no matter how low the starting rate. This creates stress on the model when the starting rate is among the lowest ever observed. It is not enough to generate scenarios that contain rates lower than observed, but to average lower over 30 years requires a significant stretch outside of calibration to historical data. While the concern over "low for long" is certainly justified, we recommend revisiting the acceptance criteria for stochastic scenarios or developing a stress test approach to the issue using one or more deterministic scenarios.

Problem 8: Conning's approach to adjusting the mean reversion point results in changing the implicit market price of risk whenever the mean reversion point changes. Such linkage is not supported by any known theory.

Conning's approach to adjusting the mean reversion point involves changing the real-world parameters but not the risk-neutral parameters. In a real-world arbitrage-free model, the market price of risk is implicit in the difference between the two sets of parameters. Changing one without changing the other has the effect of changing the implicit market price of risk. We believe the market price of risk is a fundamental assumption that should not be linked to the mean reversion point.

Conclusion

Conning's interest rate model and its calibration seem optimized for use in a risk-neutral setting as commonly applied under IFRS 17 and Solvency II. In the real-world context several problems arise that make that form of generator less than ideally suited to the task.

There are other forms of generators that avoid these problems and are more suited to use in the real-world context required under the NAIC principle-based approach.