



Understanding the Connection between Real-World and Risk- Neutral Scenario Generators

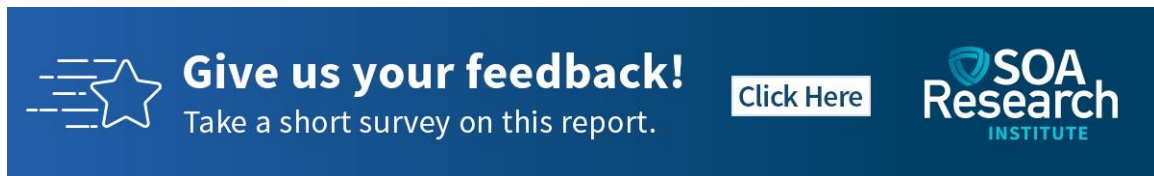
AUGUST | 2022





Understanding the Connection between Real-World and Risk-Neutral Scenario Generators

AUTHOR Stephen J. Strommen, FSA, CERA, MAAA

SPONSOR Quantitative Finance and Investment Curriculum Committee

A dark blue horizontal banner with white text and icons. On the left is a white star icon with horizontal lines extending from its left side. To the right of the star is the text "Give us your feedback!" in a bold font, followed by "Take a short survey on this report." in a smaller font. Further right is a white rectangular button with the text "Click Here" in blue. On the far right is the SOA Research Institute logo in white.

 **Give us your feedback!**
Take a short survey on this report. [Click Here](#) 

Caveat and Disclaimer

The opinions expressed and conclusions reached by the authors are their own and do not represent any official position or opinion of the Society of Actuaries Research Institute, the Society of Actuaries or its members. The Society of Actuaries Research Institute makes no representation or warranty to the accuracy of the information.

Copyright © 2022 by the Society of Actuaries Research Institute. All rights reserved.

CONTENTS

Introduction	4
Section 1: Why are there two kinds of generators? (Fitness for purpose)	5
Section 2: What it means to be market-consistent	8
2.1 Reflecting expectations	8
2.2 Common mathematical form of scenario generators	9
2.3 The common connection between real-world and risk-neutral calibrations	9
Section 3: Interest rate models	10
3.1 Market behavior in interest rate models	11
3.2 The market price of risk in interest rate models	13
Section 4: Equity index models	16
4.1 Market behavior in equity index models	16
Alternatives for the expected drift term	16
Alternatives for the random shock term	17
4.2 Market price of risk in equity index models	17
Section 5: Calibration of market-consistent and market-coherent models	18
5.1 General approach	18
5.2 Risk-neutral calibration	18
5.3 Real-world calibration	19
5.4 Combined calibration: a market-coherent real-world model	20
Section 6: Acknowledgments	23
Appendix A: The P-measure and the Q-measure	24
Interest rates	24
Equity returns	27
Appendix B: Some more complex interest rate models	28
Double mean reverting models.....	28
The three-factor Cox-Ingersoll-Ross model	30
Appendix C: The meaning of arbitrage-free	32
Appendix D: Term premiums and the market price of risk	33
Notes on calibration of term premiums	35
Appendix E: Fair value, Principle-based Reserves, and the Valuation Manual	37
Obtaining scenario-specific values	37
Blending scenario-specific values into a single reported value	38
Selecting a scenario generator that is fit for purpose	39
References	40
About The Society of Actuaries Research Institute	41

Understanding the Connection between Real-World and Risk-Neutral Scenario Generators

Introduction

Scenario generators are sometimes described as being as either “real-world” or “risk-neutral”. The difference between real-world and risk-neutral generators is in the treatment of the market price of risk. This document focuses on the difference in treatment of the market price of risk, and helps the reader understand the relationship between real-world and risk-neutral scenarios and when it is appropriate to use each type.

Risk-neutral generators are referred to as “market-consistent” because they are calibrated to reproduce market prices of traded financial instruments. It is often assumed that real-world generators cannot reproduce market prices, but in fact the ability to reproduce market prices is a matter of calibration and a real-world generator can be calibrated to reproduce market prices. We introduce the term “market-coherent” to describe a real-world generator that is calibrated to reproduce market prices, explain how such calibration can be done, and discuss why it can be useful. The term “market-coherent” was chosen because such a generator not only reproduces current market prices but also reflects the real-world evolution of those prices into the future by using the P-measure (real-world probabilities) rather than the Q-measure (risk-neutral probabilities).

To help illustrate how real-world and risk-neutral methods can be calibrated to produce the same market price, an Excel workbook has been prepared to accompany this paper. The workbook illustrates the valuation of a simple put option on equity stock using real-world and risk-neutral methods. The probability distributions associated with real-world and risk-neutral probabilities are shown side-by-side, and the means by which the methods use different probabilities to arrive at the same value are shown in detail.



Give us your feedback!

Take a short survey on this report.

[Click Here](#)

SOA
Research
INSTITUTE

Section 1: Why are there two kinds of generators? (Fitness for purpose)

Before addressing the connection between real-world and risk-neutral scenario generators, it is important to understand why there are two different kinds of generators in the first place. The two kinds of generators were created to address two fundamentally different problems. That means they are fit for different purposes, and it is important when choosing a generator to choose one that is fit for the purpose at hand.

Real-world generators were created to provide realistic stochastic simulation of the path of future economic conditions through time. Each path is a scenario, and such realistic but stochastically generated scenarios can be used in financial models to quantify risk. By running simulations using such scenarios, one can try to answer the questions “How bad could things get” and “How likely is that?” in the context of risk management for a financial institution.

Therefore, real-world scenario generators are fit for the purpose of simulation across time.

Risk-neutral generators were created for market-consistent valuation of options and other derivatives with an uncertain payoff. A market-consistent valuation method is one which reproduces the market price of market-traded instruments including options and derivatives. Not all such options and derivatives are market-traded and have observable market prices. When buying or selling a derivative that is not traded widely, it is very useful to determine its value by comparison with similar derivatives that are widely traded. It is useful to determine a “market-consistent” price or value. Risk-neutral methods, and generators based upon them, allow determination of a market-consistent value without making any assumption about the market price of risk. That is useful because the price of risk is not directly observable, but the market prices used to calibrate a risk-neutral generator are observable.


Therefore, risk-neutral scenario generators are fit for the purpose of valuing an item with an uncertain payoff in a market-consistent way.

The world would be much simpler if this distinction in purpose were always clear-cut. It is more complex when the problem at hand is to value an item with an uncertain payoff by means of stochastic simulation across time. In the realm of accounting for insurance contracts, that kind of valuation approach has been widely adopted. In that context the choice as to which kind of scenario generator is fit for the purpose requires understanding the difference between them at a deeper level.

One key to understanding this difference is to consider the universe of all possible scenarios. In abstract terms, any scenario generator is a means of picking scenarios from that universe. The parameters of the generator essentially assign a probability of selection to each possible scenario. The probability of selection differs in a fundamental and systematic way between real-world and risk-neutral scenarios. Real-world scenarios use the P-measure, or real-world probabilities. Risk-neutral scenarios use the Q-measure, or risk-neutral probabilities. These different measures assign different probabilities to the same set of events in the future.


The chart below illustrates this concept in a very simple case. Suppose the problem at hand involves projecting the future value of a share of equity stock one year from now. The universe of possible scenarios ranges from very low value to very high value. At the starting date of the projection, the real-world probabilities are for a mean expected value near the middle of that range. Real-world scenarios are probability weighted around that mean expected value.

Figure 1
DISTRIBUTION OF REAL-WORLD SCENARIOS



What about risk-neutral scenarios? The next chart adds the probability distribution for a set of risk-neutral scenarios. Note that there is a lot of overlap between the risk-neutral and real-world distributions, indicating that many of the same scenarios may appear in both sets. Note also that the risk-neutral distribution is not centered on the same central value: it is centered on a lower value. This means that the average projected value of a risky investment in risk-neutral scenarios will be lower than in real-world scenarios. The probability weighting in risk-neutral scenarios (Q-measure) gives more weight to adverse results (lower projected value in this case) than the P-measure. For example, the central value in the risk-neutral probability weighting is based on the price increasing at the risk-free rate. That rate is typically lower than a realistic market expectation that would include a risk premium.

Figure 2
DISTRIBUTION OF REAL-WORLD AND RISK-NEUTRAL SCENARIOS



In these charts the horizontal axis is a projected value. Keep in mind that the starting value is the same in both real-world and risk-neutral scenarios. And, importantly, the expected present value of the stochastic future price is the same in both the real-world and risk-neutral scenario sets. That's because the discount rate differs between real-world and risk-neutral methods. The risk-neutral discount rate is always the risk-free short-term rate. The real-world discount rate is higher because it includes a risk premium for the equity price risk. That risk premium exists because real-world markets are risk-averse; an uncertain future price is discounted at a rate that includes a risk premium related to the degree of uncertainty in the future price. Therefore, projected values and prices are higher in real-world market-coherent scenarios than in risk-neutral scenarios, but the present values are the same. This means that projections of value based on risk-neutral scenarios diverge more and more from real-world projections the further into the future you go. Risk-neutral methods are clearly meant for market-consistent valuation on the scenario start date, where they match current real-world prices, but not for projections or simulations that involve decision-making based on simulated conditions on dates in the future. In risk-neutral scenarios, simulated future conditions do not have a realistic probability distribution.


This has implications when using stochastic scenario methods for valuation of insurance contracts or pension obligations. When the valuation basis is market value, risk-neutral and real-world methods can (in theory) produce the same present value in cases where there are no future decisions to be made based on future market conditions. That would be the case for a block of insurance contracts or pensions backed by a true replicating portfolio. But in many cases simulation of obligations involves decisions to be made by either the administrator or the consumer based on future conditions¹. These decisions have financial effects that are not considered in the calibration of risk-neutral scenarios because that calibration is based on other financial instruments. Therefore, use of the risk-neutral probabilities (Q-measure) to probability-weight the financial effects of such decisions is less than ideal because those are not the real probabilities. It would be ideal if the real-world probabilities (P-measure) could be used.

Using real-world scenarios and the P-measure for this purpose requires that they be calibrated to reproduce market prices. Such a calibration requires an assumption concerning the market price of risk, which is not directly observable. Because of this, and since risk-neutral methods and the Q-measure do not require an assumption for the market price of risk, risk-neutral methods are commonly used. Nevertheless, the market price of risk can be inferred indirectly using some generally accepted models. With a reasonable assumption for the market price of risk, the P-measure can be calibrated to reproduce market prices. We will call real-world scenarios based on such a calibration "market-coherent" because they not only reproduce current market prices but also employ real-world probabilities (the P-measure) concerning the evolution of financial conditions in the future. Market-coherent scenarios can be used in valuation of the kind of obligations described above.

What about real-world scenarios that are not market-coherent? Such generators can be useful for risk management or in a regulatory context where black swan events are an important consideration. The next chart illustrates one possible distribution for such a set of scenarios. In this case the distribution is wider with longer tails and would not reproduce market prices of many derivatives. For this kind of purpose, the tail scenarios are the focus of attention, and the exact center of the distribution is less important.

¹ Such future decisions on the part of the insured may include whether to keep a contract in force because guarantees have become valuable or surrender it because guarantees no longer have much value. Decisions on the part of the insurer may include how to invest future renewal premiums. The distribution of future market prices and interest rates that would drive future investment decisions are not realistic in risk-neutral scenarios.

Figure 3
DISTRIBUTION OF REAL-WORLD SCENARIOS WITH EXTRA TAIL RISK



Section 2: What it means to be market-consistent

2.1 REFLECTING EXPECTATIONS

The term “market-consistent” means a valuation or a set of scenarios must be consistent with current market prices.

Market prices for financial instruments reflect expectations of the future. For example, the market price of fixed income investments will be lower if it is expected that interest rates will rise and higher if it is expected that interest rates will fall in the future. The market price of equity investments like corporate stock depends on expectations regarding future earnings. The market price of options depends not only on the strike price but also on the expected volatility of market prices over the term of the contract because the probability of an out-of-the money option having value depends on the volatility.

A calibration that reproduces current market prices focuses mainly on two aspects of market expectations:

- the central path or direction of future prices
- the volatility around that path; how far actual returns might vary from the expected path

Calibration of a scenario generator to reproduce current market prices involves using observed market prices to quantify the central path and volatility around that path. Real-world and risk-neutral calibration methods infer different central paths and different volatilities. If real-world and risk-neutral scenarios were used in the same way to calculate a present value, they would produce different values. But real-world valuation methods and risk-neutral valuation methods use the scenarios differently, and that difference allows them to both reproduce the same market price. An Excel workbook accompanying this paper illustrates this idea using an example valuation of a put option on equity stock.

Since both real-world and risk-neutral methods can be calibrated to reproduce market prices, under the definition above they could both be categorized as market-consistent when calibrated that way. Common usage has come to associate the term “market-consistent” only with the risk-neutral approach. To avoid confusion, the term “market-coherent” can be used to refer to a real-world generator calibrated to reproduce market prices.

Scenario generators also need to reflect understandings about how markets behave. These include the idea that markets are arbitrage-free and that market participants are risk-averse. Scenario generators can be designed and calibrated to embody those expectations as well.

There are some popular forms of generators that are not arbitrage-free. These are commonly used for real-world simulation. Since they are not arbitrage-free they do not reproduce market prices. There is generally no risk-neutral version of such models, so one cannot explain the connection between risk-neutral and real-world calibrations and they will not be discussed further in this document.

2.2 COMMON MATHEMATICAL FORM OF SCENARIO GENERATORS

The common mathematical form of a scenario generator is simple:

$$(next\ period\ value) = (current\ value) + (expected\ drift) + (random\ shock)$$

The expected drift is calibrated to the central expectation of the change in value. The random shock is calibrated to volatility in the value.

2.3 THE COMMON CONNECTION BETWEEN REAL-WORLD AND RISK-NEUTRAL CALIBRATIONS

Consider the common mathematical form shown above. Taken together, the terms for the expected drift and the random shock define a probability distribution for next period’s value. A key to understanding the connection between real-world and risk-neutral models is that the distribution in a real-world calibration is different from the distribution in a risk-neutral calibration. In particular, the drift term is always different between real-world and risk-neutral calibrations.

The risk-neutral approach to calibration is based on the idea that the market price of risk is zero and the expected drift in the price of all investments is the short-term risk-free rate. But since the real-world drift in price is higher for riskier investments, the lower expected drift in a risk-neutral calibration amounts to giving more probability weight to adverse events that result in lower future market prices. That is how a risk-neutral calibration reflects risk aversion; it effectively gives more probability weight to adverse future events.

This is the key to understanding why the intended usage of risk-neutral scenarios is limited to valuation. Risk-neutral scenarios are unrealistic in that they embed an implicit margin for risk by giving more weight to adverse events – that is, events that lead to lower market prices. But that is exactly why they are useful for valuation purposes. Valuations using risk-neutral techniques include risk margins implicitly and avoid the need to add an explicit risk margin. That is why, when using risk-neutral scenarios, the estimated value is an average across all scenarios with no add-on margin, but when using real-world scenarios some sort of risk margin needs to be included. One common means of adding a risk margin in real-world valuations is to set the value to the contingent

tail expectation (CTE), the average of the subset of scenarios containing the most adverse results rather than the average across all scenarios.²

There are countless different implementations of the common mathematical form shown above. The implementation depends on the kind of value for which scenarios are to be generated, and on the kind of market behavior over time that one wants to simulate. The following sections illustrate some generator formulations for interest rates and for a stock price index that includes dividend re-investment. Differences between various formulations are explained in terms of the market behavior they simulate and the way the market price of risk is embedded.

Section 3: Interest rate models

Arbitrage-free interest rate models focus on the path of the short-term risk-free interest rate. The remainder of the risk-free yield curve at any time is defined as an expectation based upon the path of the short-term risk-free rate. The expectation operator is applied to the discounted value of a future payment using path-wise discounting. The expected discounted values are treated as spot prices, and those spot prices define spot rates that represent the yield curve.

If the short-term risk-free rate at time t is r_t then the spot price P for maturity $t=T$ is the risk-neutral expectation³:

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

The expression for the spot price given above applies to risk-neutral calibrations. Real-world calibrations simulate movement of the short-term rate differently and therefore are not consistent with this formula for the spot price.

Why is the formula defining movement of the short-term rate different between real-world and risk-neutral models? The reason is the difference in purpose. The purpose of a real-world model is for simulation of the actual path of the short-term risk-free rate. The purpose of a risk-neutral model is valuation, so the model must produce forward paths of short-term rates that reproduce the prices of pure discount bonds based on the formula shown above. The forward path in a risk-neutral model is different from the path in a real-world model because of the risk associated with locking in the fixed interest rate in a long-term pure discount bond. That risk is reflected in the

² Another common means of adding a risk margin in real-world valuation is to start with the average value across all real-world scenarios and then add an amount that represents the present value of the cost of capital required to manage the risk of the item being valued. The theory in that approach is that the cost of capital represents the market price of risk.

³ The spot price is not the same as the discounted value using the expected path of the short-term rate, which would be expressed as follows:

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

This expression may have a very similar value, but it is not the same. Since both expressions involve an expectation, they are often confused. The situation here is somewhat analogous to measures of central tendency in a probability distribution, where the mean and the median are both measures of central tendency, but they are not the same.

volatility of the market price of a pure discount bond – volatility that does not exist in the price of a short-term money market fund earning the short-term risk-free rate. Investors in pure discount bonds demand a higher return as a reward for taking that risk, so the forward rate paths consistent with the price of a pure discount bond are higher than the forward paths of the truly risk-free short-term rate.

Previously it was said that risk-neutral scenarios give more weight to adverse events, meaning declines in prices. Yet now we say that forward rate paths in a risk-neutral model are higher than in a real-world projection. How is this consistent? It is consistent because an upward movement in the forward rate leads to a decline in the price. Higher interest rates mean lower prices. Giving more weight to adverse events that decrease the market price means giving more weight to increases in interest rates that decrease the market price.

Risk-neutral scenarios are not intended for use in simulation where those future higher interest rates might lead to higher returns on future new investments. Risk-neutral scenarios are intended for use in valuation of investments purchased in the past. Any use of risk-neutral scenarios to simulate investment purchases in the future violates the intended purpose of such scenarios for these reasons⁴:

- It results in purchase of future investments at higher yields (i.e., more favorable future conditions) than the market actually expects.
- It is inconsistent with the financial theory underlying risk-neutral models which was developed for the purpose of valuation of existing investments, not simulation of future investments.

Given a formulation for r_t it is often possible to derive a closed-form expression for the spot price consistent with the formula above involving the central expectation along possible paths. A closed-form expression is desirable because it allows the shape of the full yield curve to be calculated directly. When such a closed-form expression does not exist, spot prices and the yield curve are approximated using Monte Carlo simulation of many scenarios for r_t to estimate the central expectation in the formula above.

3.1 MARKET BEHAVIOR IN INTEREST RATE MODELS

There are two important and separate characteristics to any arbitrage-free model for interest rates, and both center on the formula defining movements of the short-term risk-free rate. One characteristic is the kind of general market behavior that is simulated. The other characteristic is whether the calibration is real-world or risk-neutral in nature.

In what follows we first discuss a few different models of general market behavior. Then we discuss the difference between real-world and risk-neutral versions of each model. This is done to emphasize that when choosing an ESG, one needs to be aware of both characteristics of the model. All market-consistent models are not alike and do not produce identical results, and the differences are not limited to whether they are real-world or risk-neutral; the differences depend also on the underlying model of market behavior.

Perhaps the best-known and simplest interest rate model is the Vasicek model. The Vasicek model treats the short-term risk-free interest rate as a mean-reverting random walk. Using the common mathematical form from above, the Vasicek model in discrete time is expressed as follows:

⁴ The underlying theory does not support this usage in an asset-liability simulation where there is re-investment risk or where there are options in the contracts that create decision points where the decision will be based on future economic conditions. In those situations, use of real-world market-coherent scenarios is more consistent with the underlying theory.

(next period value) = (current value) + (expected drift) + (random shock)

$$r_{t+1} = r_t + \kappa(\theta - r_t) + \sigma W_t$$

κ = mean reversion strength or speed

θ = mean reversion point

σ = volatility

W_t = a random number from a Gaussian distribution with mean zero and variance one

To develop a closed-form expression for the spot price and the remainder of the yield curve, the discrete model above is expressed as a stochastic differential equation in continuous time⁵:

$$dr = \kappa(\theta - r)dt + \sigma dW$$

The Vasicek model assumes a certain kind of market behavior. Aspects of that behavior are that the volatility is constant and that interest rates can become negative if enough negative random shocks accumulate.

Other models have been developed assuming different kinds of market behavior. For example, one can observe that volatility is not constant but tends to be higher when interest rates are higher. One can also observe that there is a strong tendency for interest rates not to go below zero. The Cox-Ingersoll-Ross model and the Black-Karasinski model each embed those behaviors in different ways.

Here is the Cox-Ingersoll-Ross model:

$$dr = \kappa(\theta - r)dt + \sigma\sqrt{r}dW$$

The only difference from the Vasicek model is the random shock term which is multiplied by the square root of the current rate. This makes the random shocks (i.e. the volatility) depend on the current level of the interest rate. When interest rates approach zero, the random shocks get very small and mean reversion forces the interest rate back up toward the mean reversion point. This prevents interest rates from ever getting to or below zero.

Here is the Black-Karasinski model:

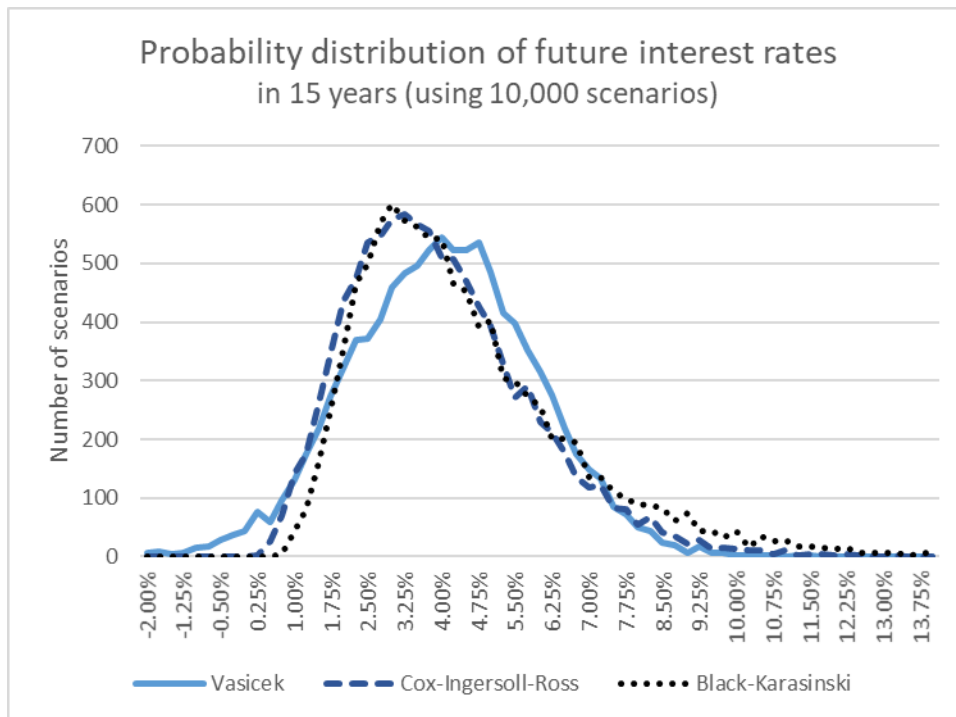
$$d \ln(r) = \kappa(\ln(\theta) - \ln(r))dt + \sigma dW$$

This model is also a mean-reverting random walk, but the simulation is done with the logarithm of the interest rate rather than the interest rate itself. The volatility term is constant in absolute terms, but in proportion to the logarithm of the interest rate the volatility is smaller when interest rates are lower and vice versa, as in the Cox-Ingersoll-Ross model. But when the interest rate approaches zero, the magnitude of its logarithm approaches infinity and that makes the mean reversion term very sensitive to the level of interest rates. In this model the mean reversion term explodes when interest approach zero, providing an even stronger reversion to the mean than in the Cox-Ingersoll-Ross model. And when interest rates are high, the mean reversion term gets substantially smaller and has less effect.

⁵ Note that the values of the parameters to be used in the discrete time version of the model are not the same as those in the continuous time version. The values depend on the size of the discrete time step in use.

These differences in assumed market behavior are illustrated in the chart below. The chart shows the distribution of the short-term interest rate 15 years in the future based on each of the models, using comparable calibrations. The distributions were simulated using 10,000 scenarios and counting the number of scenarios that fell into each 0.25% bucket. The distribution for the Vasicek model is a bell-shaped curve that includes some negative interest rates. The distribution of the Cox-Ingersoll-Ross model is skewed with a longer tail on the high end and does not include any negative interest rates. The distribution for the Black-Karasinski model is even more skewed.

Figure 4
SIMULATED INTEREST RATE DISTRIBUTIONS FROM THREE MODELS



All three of these distributions are based on calibrations with the same mean reversion point, the same value for the mean reversion strength parameter, and the same volatility at the mean. The point of showing the difference between them is to motivate the idea that the choice of model affects the valuation because it reflects the model's assumptions about market behavior and the likelihood of different events in the future. This dependence of any market-consistent valuation on the choice of underlying model should not be overlooked.

3.2 THE MARKET PRICE OF RISK IN INTEREST RATE MODELS

The equations above do not explicitly include the market price of risk, which is normally denoted as λ . Let's briefly review what the market price of risk is and then see how it has been incorporated into these models.

To define the market price of risk we start with the idea that risk-averse investors require a higher expected or average return on riskier investments. We measure risk by the volatility of the market price. Greater volatility and greater risk increase the probability that the price may decline causing a loss, but they are associated with higher returns on the average. With that in mind we define the market price of risk as a ratio:

$$\text{Market price of risk} = \lambda = \frac{(\text{Expected average return}) - r}{\text{standard deviation of price}}$$

The excess average return over the risk-free rate is called a risk premium. The market price of risk is the ratio of the risk premium to the volatility of the price.

One might ask why the market price of risk should be incorporated into a formula defining movements of the short-term risk-free rate. After all, since the rate is risk-free, then the market price of risk does not apply. The explanation is the formula given earlier for the spot price of a pure discount bond as an expectation based on the future movement of the short-term rate. Pure discount bonds are not risk-free; their price is volatile because the market interest rates used to discount their fixed future payoff are volatile. Since the price of a pure discount bond is volatile, the market expects a return higher than the risk-free rate and uses that higher return to calculate the market price. The path of the short-term rate used to price pure discount bonds (the risk-neutral path) must be higher than the real-world expected path to a degree that depends on the market price of risk.

Therefore, it is common to introduce the market price of risk into these formulas as an upward drift in the risk-free rate, an upward drift equal to the volatility times λ . For example, the Vasicek formula including λ is this⁶:

$$dr = \kappa(\theta - r)dt + \sigma(dW + \lambda)$$

Note that this is equivalent to simply increasing the mean reversion point parameter:

$$dr = \kappa(\theta - r)dt + \sigma\lambda + \sigma dW$$

$$dr = \kappa\left(\left(\theta + \frac{\sigma\lambda}{\kappa}\right) - r\right)dt + \sigma dW$$

The last expression above is identical to the original expression for the Vasicek model, but with the mean reversion parameter θ replaced by $\left(\theta + \frac{\sigma\lambda}{\kappa}\right)$. Because of that, sometimes the market price of risk in a model like this is re-defined to be the value added to the mean reversion parameter rather than the more abstract concept of risk premium divided by the volatility.

The simple relationships shown above help one understand some of the language commonly used with reference to these models.

The P-measure and the Q-measure. A measure refers to the probability distribution of future movements of the variable under consideration, in this case dr . As shown above, in an arbitrage-free model the spot price is defined (measured) in terms of an expectation based on the path of r . Both the drift term and the random shock term are part of the measure. As was shown above, the change in measure from real-world (the P-measure) to risk-neutral (the Q-measure) in the Vasicek model is accomplished by modifying the random shock term, replacing dW with $d\widehat{W} = dW + \lambda$. Equivalently, the change in measure can be accomplished by changing the drift term, replacing θ with $\widehat{\theta} = \theta + \frac{\sigma\lambda}{\kappa}$. This equivalence between changing the random term and the drift term is what makes this and many similar models mathematically tractable so that a fixed-form equation for the spot price at any maturity can be derived. This equivalence is also what makes the P-measure and the Q-measure “equivalent Martingale measures”. An appendix is included to provide further illustrations and explanations of the relationship between the P-measure and the Q-measure.

⁶ In many references, lambda is introduced with a negative sign rather than a positive sign. When that is done, calibration shows that lambda has a negative value. Here we introduce it with a positive sign and a positive value because the market price of risk is defined in a way that anticipates a positive value.

The real-world parameters vs. the risk-neutral parameters. This expression refers to the values of the parameters, not their definitions. The mathematical definition of the parameters is the same in real-world and risk-neutral versions of the model. The parameters of the Vasicek model are κ , θ , and σ . The difference between real-world and risk-neutral parameter values is in the value of θ . If the real-world value of that parameter is $\theta = x$ then the risk-neutral value of that parameter is $\hat{\theta} = x + \frac{\sigma\lambda}{\kappa}$ as shown in the formulas above⁷.

Let's turn now to the Cox-Ingersoll-Ross model. The market price of risk can be introduced into the Cox-Ingersoll-Ross model in at least two different ways. The difference illustrates the complexity of reflecting the market price of risk in other models.

The first method is to replace the random term dW in the Cox-Ingersoll-Ross model with $d\tilde{W}_t = dW_t - \frac{\lambda\sqrt{r}}{\sigma} dt$. This is the approach used in the original Cox-Ingersoll-Ross paper⁸.

$$\begin{aligned} dr_t &= \kappa(\theta - r_t)dt + \sigma\sqrt{r_t}d\tilde{W}_t \\ dr_t &= \kappa(\theta - r_t)dt + \sigma\sqrt{r_t}(dW_t - \frac{\lambda\sqrt{r}}{\sigma} dt) \\ dr_t &= \kappa(\theta - r_t)dt - \lambda r_t dt + \sigma\sqrt{r_t}dW_t \\ dr_t &= (\kappa + \lambda) \left(\frac{\kappa\theta}{(\kappa + \lambda)} - r_t \right) dt + \sigma\sqrt{r_t}dW_t \end{aligned}$$

In the last line above we see that the terms containing lambda have been removed from the random term and put equivalently into the drift term. The parameter κ has been replaced by $\hat{\kappa} = \kappa + \lambda$ and the parameter θ has been replaced by $\hat{\theta} = \frac{\kappa\theta}{\kappa + \lambda}$. One can use this form of the model to derive a formulaic expression for the spot price. The formulaic expression is very complex and is shown in *Economic Scenario Generators – A Practical Guide* (Pedersen et. al, 2016).

Another method for introducing the market price of risk is to notice the following:

$$dr_t = \kappa(\theta - r_t)dt + \sigma\sqrt{r_t}d\tilde{W}_t = \hat{\kappa}(\hat{\theta} - r_t)dt + \sigma\sqrt{r_t}dW_t$$

The second form of the model at the end of the line above can be used to develop a formulaic expression for the spot price that depends on adjusted values of κ and θ without any reference to lambda.

If the adjustments are γ_1 and γ_2 then define $\hat{\kappa} = \kappa + \gamma_1$ and $\hat{\theta} = \theta + \gamma_2$. The effect of the two adjustments together is assumed to reflect the market price of risk. The value of this approach is that if one is only interested in risk-neutral valuation, then one can use the formula for the spot price based on $\hat{\kappa}$ and $\hat{\theta}$ to calibrate their values

⁷ One might be confused by the fact that lambda appears in this formula for the risk-neutral parameter value $\hat{\theta}$. Isn't the market price of risk zero in a risk-neutral calibration? The explanation is that calibration can be done in a way that estimates the value of $\hat{\theta}$ directly without the need to adjust for the market price of risk. The market price of risk is not zero, rather it is reflected implicitly so no adjustment for the market price of risk is required. A risk-neutral calibration does not really assume the market price of risk is zero; it could not reproduce market prices if it did. Rather, a risk-neutral calibration provides parameter values that reflect the market price of investment risk implicitly, so no explicit adjustment for investment risk is required in a risk-neutral valuation.

⁸ This is one of the references in which lambda has a negative value.

directly to market prices and never need to specify the implicit values for γ_1 and γ_2 or the real-world values κ and θ .

The simple interest rate models discussed above are not commonly used in practice because their simplicity does not reproduce actual market behavior very well. An appendix is included in this document to illustrate the treatment of the market price of risk in some more complex interest rate models.

The market price of risk can seem obscure when presented as a factor in a mathematical formula. Understanding can be improved by viewing the effect of the market price of risk on the yield curve. Recall that there is often a closed-form formula for the yield curve based on a definition of the process for r_t . The result of the formula depends on the parameter values used. When there are both real-world and risk-neutral parameter values for the same model, they will produce different yield curves starting from the same short-term rate. The difference between those yield curves is caused by the market price of risk because the market price of risk is the difference in the parameters. The difference between those yield curves is a set of “term premiums” that vary by term to maturity.

Term premiums arise because of the risk associated with locking in a fixed rate for a long time in a world where future interest rates are uncertain. Term premiums represent the market price of that risk. One can judge the reasonableness of any formula for the market price of risk based on the term premiums that it implies, that is, based on the difference between yield curves produced using real-world and risk-neutral parameter values.

Section 4: Equity index models

Arbitrage-free equity index models focus on the total rate of return of investments in the equity index. Here we consider only models that include re-investment of dividends. Such models follow the general form outlined above:

$$(next\ period\ value) = (current\ value) + (expected\ drift) + (random\ shock)$$

Generally, these models are simpler than interest rate models. The expected drift for any period is typically the sum of the short-term risk-free rate and a risk premium. The realized returns in any specific scenario are dominated by the random shocks which tend to be much larger in magnitude than the expected drift.

In risk-neutral calibrations of these models, the risk premium is zero so the expected drift is always equal to the short-term risk-free interest rate. Differences between risk-neutral versions of such models are limited to the random shock term.

4.1 MARKET BEHAVIOR IN EQUITY INDEX MODELS

The design of an equity index model focuses on ways to model the expected drift and the random shock (volatility). When multiple equity indices are included in one ESG, there is focus on the relationships of the drift and volatility of various indices. We will not address relationships between separate equity indices in this document.

ALTERNATIVES FOR THE EXPECTED DRIFT TERM

There are no alternatives for behavior of the expected drift in a risk-neutral model. The expected drift in a risk-neutral model is always the current short-term risk-free rate.

Alternatives for market behavior of the expected drift in a real-world model include:

- Whether the expected drift is the sum of the current short-term risk-free rate and a fixed risk premium or is constant. If the expected drift is a constant, then the implied risk premium changes over time to offset changes in the short-term risk-free rate.
- Whether the expected drift is regime-switching.
- Whether the expected drift is mean-reverting over time, so that the level of the equity index tends to mean-revert to a level that grows at a more stable compounded rate over long periods.

ALTERNATIVES FOR THE RANDOM SHOCK TERM

Alternatives for market behavior of the random shock include:

- Whether the volatility of random shocks is constant or is stochastic or regime-switching in nature.
- Whether the volatility of random shocks is supplemented by a “jump process” that creates large jumps in value at random intervals.

Many of these alternatives are discussed in *Economic Scenario Generators – A Practical Guide* (Pedersen et. al. 2016) and are not elaborated here.

4.2 MARKET PRICE OF RISK IN EQUITY INDEX MODELS

The market price of risk in an equity index model is reflected in the excess of the drift term over the short-term risk-free rate. That excess is the equity risk premium. In a risk-neutral model the drift term is equal to the short-term risk-free rate so the market price of risk and the equity risk premium are zero.

Risk-neutral equity models are intended for valuation of options and other derivatives. One does not need a model to perform valuation of a stock index – the value is the current level of the index. Since the drift term in a risk-neutral equity model is always the short-term risk-free rate, the valuation of options depends only on the random shock term in the model. Risk-neutral models can vary in the way the random shock term is constructed, and calibration is important. The basic risk-neutral equity model is the Black-Scholes model in which volatility is Gaussian and calibration amounts to determining the implied volatility level that corresponds to an observed market price.

Implied volatility in the Black-Scholes model is determined by calibration to the market price of an option. In the Black-Scholes model the drift in price is the short-term risk-free rate and the only variable is the volatility. Implied volatility is the volatility which, when used in the Black-Scholes formula, leads to an option value that equals the market price of that option. When this is done, implied volatility tends to depend on the strike price and tenor of the option used for calibration. The variation by strike price and tenor is called the “volatility surface”. The volatility surface is very useful for valuation. But it presents a problem when the purpose of a simulation is not valuation of an option at a point in time but focuses on the future path of the underlying index. In that context, a choice must be made concerning the level of volatility for the underlying index in the scenarios. Reasonable choices come from the range of volatilities in the volatility surface.⁹

⁹ The volatility surface corresponds to the Black-Scholes model, which is a constant volatility model. More complex models involve stochastic volatility, and such models are often calibrated to value options using simulation. In that context, the parameters of a stochastic volatility model are calibrated to market prices of some options, and the result may be several sets of parameters that differ by the strike and tenor of the option. An analogous issue arises there because when the purpose of a simulation is not valuation of an option at a point in time but focuses on the future path of the underlying index, a single set of parameters is needed to generate scenarios for the path of the index.

In a real-world model the equity risk premium tends to be positive. But some of the options for real-world models involve the idea that equity returns are driven by forces other than a risk premium over the risk-free rate. Such forces may include fiscal and monetary policy, GDP growth, unemployment, and others. Because of those other influences, the equity risk premium may not be constant and the drift in equity prices may not be tied to the sum of the risk-free rate and a fixed risk premium.

Section 5: Calibration of market-consistent and market-coherent models

5.1 GENERAL APPROACH

Economic Scenario Generators – A Practical Guide (Pedersen et. al. 2016) provides a good discussion of calibration procedures in section 7. Importantly, it defines separate calibration approaches for risk-neutral and real-world calibrations. Here we generalize so that the two calibration approaches can be viewed as two versions of the same process. The generalized process involves iterative adjustment of parameter values to improve fit to historical observations. The difference between risk-neutral and real-world calibration is not in the use of iterative fitting, it is in selecting which historical observations are used as targets of the ESG's fit, and how the goodness of fit is measured.

No model yet invented can perfectly fit all historical observations. There will always be some historical data that a model can fit better than other data. Historical data that is not used directly in calibrating the model may not be fit well at all. As noted above, risk-neutral and real-world calibrations focus on fitting different aspects of historical data. Because of this, it should not come as a surprise that sometimes risk-neutral and real-world calibrations are very different from one another and appear inconsistent. The kind of conceptual connection between risk-neutral and real-world calibrations that was presented earlier in this document can be reflected when both kinds of model are calibrated together. But often they are calibrated independently using different samples of historical data. Odd relationships between real-world and risk-neutral parameters can arise from independent calibration to separate data sets.

In the sections below we briefly discuss risk-neutral and real-world calibration separately and then explain the kind of combined calibration required for a real-world market-coherent model.

5.2 RISK-NEUTRAL CALIBRATION

The purpose of a risk-neutral ESG is valuation. Calibration of a risk-neutral ESG is the process of choosing parameters that reproduce market prices, so that the ESG can be used to determine a market-consistent value for items that are not widely traded and therefore have no observable market price.

Typically a risk-neutral ESG is formulated in a way that leads to closed-form expressions for the prices of various kinds of items traded in the market, including both direct investments and derivatives. The process of calibration amounts to iteratively adjusting the values of the parameters in the closed-form expressions for market prices to achieve the best fit to the widest array of observed market prices.

These considerations involved in risk-neutral calibration deserve special attention:

- Market conditions are constantly changing. Often (but not always) a risk-neutral ESG is re-calibrated on every valuation date to current prices in order to be most consistent with current market prices. Sometimes recalibration is less frequent, and the calibration is intended to be consistent with market prices over a period of time, but the period of time is selected to represent relatively stable market conditions.

The point here is that the time frame of the historical data sample used for risk-neutral calibration is typically shorter than that used for real-world calibration.

- Once a formula for a market price based on a set of parameters is developed, the process of fitting parameter values to reproduce market prices often does not associate any meaning to those parameters. They are just numbers to be adjusted in a mathematical optimization. The resulting parameter values may not make much sense in the real-world. This is not viewed as a problem in the context of valuation. But care must be exercised when using such risk-neutral parameter values as a starting point when building a real-world model.
- Sometimes the calibration process includes solving for different parameter values for different derivatives of the same investment. This is most common in connection with the volatility parameter. For example, in the Black-Scholes model the “volatility surface” is an array of different values for the volatility parameter depending on the strike prices and tenors of options on some underlying index. By using an array of values like this, market prices can be matched more exactly. But the user should understand that there is only one underlying item with only one underlying volatility, and that the need for an array of different volatilities for derivatives is evidence that the model does not fit the underlying item and its derivatives with the same parameters. The use of a “volatility surface” is just a way to improve the fit of the model to current prices. It is useful in the context of valuation, but not in the context of simulation because there can be only one volatility for the underlying item. In a simulation the volatility of the underlying item governs payoffs of any derivatives of that item.

5.3 REAL-WORLD CALIBRATION

Calibration of real-world models is covered well in *Economic Scenario Generators – A Practical Guide* (Pedersen et. al. 2016).

There are some key differences from risk-neutral calibration.

- The “goodness of fit” is measured differently. Instead of measuring the fit of formulaic prices to actual market prices, “goodness of fit” is measured with reference to expected movements over time in the generated scenarios. The *Practical Guide* says:
 - “Because real-world parameterizations are forward looking, they require explicit views as to how the economy will develop in the future, so they require a significant amount of expert judgment to determine the veracity of the scenarios that result from the parameterization process. In practice, real-world calibrations often are parameterized to be consistent with historical dynamics of economic variables, although the long-term steady-state levels associated with these parameterizations can differ from long-term historical averages in favor of current consensus expectations or even individual viewpoints.
- Because goodness of fit is measured in terms of movements over time, parameters like the mean reversion point and volatility have real meaning and the steady-state levels mentioned in the *Practical Guide* have real meaning. This contrasts with risk-neutral calibration where, as noted above, sometimes the parameter values diverge from their real meaning.
- The market price of risk is always non-zero in a real-world model and is always zero in a risk-neutral model. That means that the expected return of any investment that involves risk will be different on average when comparing real-world model and risk-neutral models. That makes real-world and risk-neutral calibrations of the same model different from one another.

5.4 COMBINED CALIBRATION: A MARKET-COHERENT REAL-WORLD MODEL

A real-world model can be calibrated to reproduce market prices, but this is achieved differently in the real-world context than in the risk-neutral context.

In both cases, reproducing market prices means

- Choosing a probability measure to use in calibration. The real-world measure is called the P-measure. The risk-neutral measure is called the Q-measure.

Calibrating the path of future interest rates using observed spot prices of pure discount bonds.

- Reflect market volatility around that path through calibration to some observed data.

Real-world and risk-neutral calibrations accomplish these things differently.

- Choosing a probability measure
 - Risk-neutral calibration uses the Q-measure
 - Real-world calibration uses the P-measure
- Reflecting expectations for the path of interest rates
 - Risk-neutral calibration centers the path of the short-term risk-free rate on the forward rate path implied by observed spot prices on pure discount bonds.
 - Real-world calibration centers the path of the short-term risk-free rate on the forward rate path implied by observed spot prices on pure discount bonds, adjusted downward to remove the term premiums included in observed long term bond rates for the risk of locking in a long-term rate.
- Reflecting market volatility
 - Risk-neutral calibration solves for the “implied volatility”, the volatility parameter value required in a closed-form formula for the price of a security to reproduce the market price. This often results in a “volatility surface” of different implied volatilities for different options on the same underlying security.
 - Real-world calibration uses historical realized volatility in observed market prices for the underlying security.

Before discussing further details of the real-world calibration process, it is important to emphasize that real-world market-coherent scenarios are designed to be used differently than risk-neutral market-consistent scenarios in any stochastic valuation exercise. When using risk-neutral scenarios, the valuation result is the average of the scenario-specific valuations. When using real-world scenarios, the average of the scenario-specific values needs to be adjusted by adding a margin for investment risk, either at the scenario level or in total.

There is some debate over whether the margin for risk can be determined in a reliable way. If not, then the real-world valuation cannot reproduce market prices even if it is based on scenarios calibrated in the real-world market-

coherent manner described above. In this document we take the view that risk margins can be determined in a reliable way¹⁰. Techniques for calculating risk margins are outside the scope of this document.

We will first discuss characteristics of a real-world model that allows these things to be done, and then provide more detail on how they are done.

Characteristics of a real-world market-coherent model include:

- It must start with an **arbitrage-free model framework** because a model must be arbitrage-free to be market-coherent.
- An arbitrage-free model framework can have both real-world and risk-neutral calibrations. A market-coherent real-world model must have **both real-world and risk-neutral calibrations** because both are used in scenario generation.
- The difference between the parameter values in real-world and risk-neutral calibrations represents a **specific assumption about the market price of risk**.
- The model must have a **clear division between “state variables” and “parameters”**. There are two sets of parameters (real-world and risk-neutral) but only one set of “state variables”.

The statement that such a model must have both real-world and risk-neutral calibrations deserves explanation. Both calibrations are required because they are both used when generating interest rate scenarios.

Interest rate scenarios in a market-coherent real-world model are generated using the following general process:

1. The initial values of the state variables are fit to the observed starting yield curve (spot prices) using the **risk-neutral parameters**.
2. Future scenario paths for the short-term rate are generated using the **real-world parameters** to create paths for the state variables.
3. The yield curve at any time step in a scenario path is calculated using the state variables and the **risk-neutral parameters**. This process assumes we have a model that has a closed-form formula or other fast procedure to calculate the yield curve based on the state variables and the parameters.

This is the same as the process used by a risk-neutral model, except for step 2 where the real-world parameters are used in place of the risk-neutral parameters for simulating the forward path of the state variables. Steps 1 and 3 are the same as in a risk-neutral model because those steps deal with fitting the state variables to a full yield curve at a point in time. Step 2 is different because it deals with generating a path through time based on the probabilities of different paths. In step 2 a real-world generator uses the real probabilities (P-measure) while a risk-neutral generator uses the risk-neutral probabilities (Q-measure) which implicitly give more weight to adverse events that lead to lower market prices.

¹⁰ Valuation of insurance contracts for accounting purposes requires inclusion of margins for both investment risks and insurance risks. Risk-neutral scenarios are often used to provide an implicit market-consistent margin for investment risks. But the margin for insurance risks must then be added separately. If it is possible to determine a separate margin for insurance risks, one may argue that it is also possible to determine a separate margin for investment risks, or for the sum of all risks in a block of insurance contracts.

To say this another way, the future real-world path of the short-term rate is governed by the real-world parameter values, while the forward rates that define the yield curve (and spot prices) are governed by the risk-neutral parameter values.

Equity return scenarios in a market-coherent real-world model are generated using the following general process:

1. Generate the risk-free rate for the next time step
2. Add a risk premium (expected equity return spread) based on the market price of risk
3. Add a random shock based on the volatility process

This differs from the process in a risk-neutral equity return generator because the risk premium in step 2 is zero in a risk-neutral generator. Also, the volatility in step 3 may differ.

In the modeling of interest rates, flexibility in simulating market behavior generally means using a multi-factor model for interest rates. Each factor corresponds to a “state variable” whose value represents the current state of that factor. In a one-factor interest rate model, the state variable is typically the current short-term risk-free interest rate. Multi-factor models are discussed here in the Appendix on “More complex interest rate models”.

In a multi-factor model the values of the state variables are shared between real-world and risk-neutral versions of the model, while the values of the parameters differ. The state variables define conditions on the starting date and any simulated date thereafter, while the parameters govern the stochastic paths of the state variables. The parameters remain fixed over time, but the state variables move over time.

It is this division between “state variables” and parameters that enables the creation of a market-coherent real-world model. The state variables are shared between real-world and risk-neutral versions of the model. Risk-neutral parameters are calibrated for consistency with market prices, including the spot prices in the yield curve. Real-world parameters are calibrated for consistency with movements of the state variables across points in time. In a real-world market-coherent model, the prices (yield curve) at any point in time are based on the risk-neutral parameters while the movements of the state variables are governed by the real-world parameters, as is most consistent with the way each set of parameters is calibrated.

The use of both real-world and risk-neutral parameters in the same model imposes a calibration constraint that does not exist when such parameters are used in separate models. The constraint arises because the difference between the real-world and risk-neutral parameters is the market price of risk. For the two sets of parameters to make sense together, the market price of risk that connects them must be reasonable. That means, at a minimum, that it should generally be positive so that the generated scenarios tend to show higher expected average returns for riskier investments.

Calibration of risk-neutral parameters is typically done by measuring “goodness of fit” to market prices. Calibration of real-world parameters is typically done by measuring “goodness of fit” to movements in state variables. The connection in theory is the market price of risk, but when the two calibrations are done independently, the market price of risk implied by the difference between them may not make sense. That’s because, as noted earlier, models do not fit perfectly, and the approach used in risk-neutral calibration can lead to parameter values that lose their real meaning.

Therefore, the calibration of the risk-neutral parameters in a real-world market-coherent model may be done differently than for other risk-neutral models. The real-world parameters are calibrated first and taken as given. The risk-neutral parameters are treated as the sum of the real-world parameters and adjustments that reflect the market price of risk, and those adjustments are the subject of the risk-neutral calibration.

One way to calibrate the market price of risk in an interest rate model is to set it so that it produces reasonable term premiums. As noted earlier, term premiums implied in the yield curve are defined by the difference in two yield curves. One curve is generated using the real-world parameters and the other is generated using the risk-neutral parameters which are the real-world parameters adjusted for the market price of risk. The difference between those curves is a set of term premiums by term to maturity.

The two yield curves that are compared to get the term premiums are both based on the same values of the initial state variables. Since state variables move over time, one should determine the term premiums implied starting from a range of different values for the state variables.

The process just described for calibrating the market price of risk involves a fair amount of judgment along with knowledge of what a reasonable set of term premiums should be. Generally they should be positive if one believes that markets are risk-averse, but historical studies have provided more refined estimates. An appendix to this document discusses term premiums and the market price of risk in more detail.

Section 6: Acknowledgments

The researchers' deepest gratitude goes to those without whose efforts this project could not have come to fruition: the Project Oversight Group for their diligent work overseeing, reviewing and editing this report for accuracy and relevance.

Project Oversight Group members:

Charles V Ford, FSA, MAAA, CFA

Gary Hatfield, FSA, CERA, CFA, PhD

Jason Kehrberg, FSA, MAAA, PhD

James Kosinski, FSA, MAAA, CFA, PhD

Steve Marco, ASA, CERA

Hal Pedersen, ASA, MAAA, PhD

At the Society of Actuaries Research Institute:

Doug Chandler, FSA, FCIA, Canadian retirement research actuary

Dale Hall, FSA, CERA, CFA, MAAA, SOA managing director of research


Appendix A: The P-measure and the Q-measure

This appendix provides graphic depictions of the future distributions of returns created by scenario generators, focusing on the difference between real-world and risk-neutral versions. The real-world distributions are called the P-measure and the risk-neutral distributions are called the Q-measure.

INTEREST RATES


The chart below illustrates the drift of the future short-term risk-free rate in one calibration of a Cox-Ingersoll-Ross model. Both the real-world (RW) and risk-neutral (RN) drifts are shown. Since the drift is stochastic in nature, three percentile points in the probability distribution are shown: the 10th percentile, the 50th percentile and the 90th percentile. In the scenarios charted here, the starting level of the short-term risk-free rate is 4.00% and the calibrated real-world mean reversion point is also 4.00%.

The point of this chart is to show that the risk-neutral drift is higher than the real-world drift.




The next chart shows the probability density of the short-term risk-free rate based on the same scenarios as the previous chart. The density is charted based on the number of scenarios in each 0.25% interval out of a set of 10,000 scenarios.

The point of this chart is that the risk-neutral probabilities (Q-measure) define a distribution that has a higher mean and is slightly wider than the real-world probabilities (P-measure).




Risk-neutral scenarios have been described as weighting adverse events more heavily and thereby leading to lower future values of current investments. Consider how that is reflected in a simulation of a pure discount bond that matures for \$1,000. We know the starting value, and we know the maturity value, and those must be the same in both real-world and risk-neutral scenarios. But what about the value between those two points?


Consider a situation where the current risk-free short-term rate is at its mean reversion point at the start of a simulation. In the real world the short-term rate is equally likely to move up or down, so let's look at a scenario where it remains constant, and the rest of the yield curve remains constant. The chart below shows the yield curve we'll use for illustrative purposes. The one-year spot rate is 4% and the 10-year spot rate is 6%. The 10-year rate includes term premiums that compensate for the risk of locking in the rate for 10 years.



If we generate both real-world and risk-neutral scenarios with this starting yield curve, we know that the scenarios will differ, with the risk-neutral scenario drifting higher. While the short-term rate in the real-world scenario will be unchanging, the short-term rate in the risk-neutral scenario will follow the forward rate path in the initial yield curve. This will lead to higher interest rates in the risk-neutral scenario. That means lower discounted present values for the bond in the risk-neutral scenario. The chart below illustrates the average path of the value of the bond. At every point between the start date and the maturity date, the value is lower in the risk-neutral scenario.




One can calculate the return on the bond for each projection year in each of these scenarios. In the real-world scenario the return “walks down the curve” in the usual way, showing a return equal to long-term forward rates when the maturity date is far in the future, and declining over time to the short-term rate just before maturity. In the risk-neutral scenario, all investments are assumed to earn the same short-term risk-free rate (on average), but the short-term risk-free rate follows the path of the forward rate curve. The forward rate curve starts with the short-term rate and ends with the long-term rate. Basically, the forward rate curve is earned in reverse order.




EQUITY RETURNS

For equities, the real-world probabilities define a distribution of future values that is higher than the distribution based on risk-neutral probabilities. The higher values include the market reward for taking risk. The chart below is based on simulation of 10,000 scenarios, counting the number of accumulated values in each \$0.05 interval.



When percentile points on these distributions are tracked over time one can see how the difference in probabilities affects the drift of accumulated value over time.



Appendix B: Some more complex interest rate models

The interest rate models discussed in the body of this document are not used widely in practice for simulation. They were used to explain concepts in a simple context.

The reason those simple models are not used in practice is that the market behavior that they capture is unrealistically simple. They are one-factor models where the only factor describing the current state of the economy is the short-term risk-free interest rate. Since those models are mean reverting with a fixed mean reversion point, the market behavior they imply is:

- if the short-term rate is above the mean reversion point then forward rates will trend down
- if the short-term rate is below the mean reversion point then forward rates will trend up

The only shape of yield curve such models can produce is monotonic, either upward sloping or downward sloping, depending on where the current short-term rate is relative to the mean reversion point.

Interest rate models for simulation in practice are multi-factor in nature. Such models can simulate behavior where interest rates will trend in one direction in the short run and in another direction in the long run. That makes it possible to simulate behavior consistent with the shape of market yield curves that are not monotonic; hump-shaped and valley-shaped yield curves are possible.

The role of “state variables” is a key to such models. Each “factor” in a one-factor or multi-factor model has a state variable that defines the current state of that factor. In a one-factor model, the state variable is typically the current level of the short-term risk-free interest rate. In multi-factor models, the additional factors have state variables that correspond to the current level of other variables that are included in the model.

An important aspect of multi-factor models is that each factor has both a state variable that defines its current value and parameters whose values define future stochastic paths of that state variable. The general idea is that the parameter values are anticipated to be stable over time and the movements of the state variables define the behavior of the yield curve over time. Fitting a current yield curve to such a model means selecting values for the state variables that best fit the model yield curve to the observed yield curve.

Here we describe two very different approaches to multi-factor models. In a “double mean reverting” model, each factor governs a different aspect of market behavior. In a “three-factor Cox-Ingersoll-Ross” model, each factor is a separate CIR model of the short-term risk-free rate that has its own parameter values, and the simulated short-term rate is the sum of the three factors.

DOUBLE MEAN REVERTING MODELS

A double mean reverting model has two state variables: the short-term rate and a short-term mean. The short-term rate is attracted to the short-term mean in a random walk, and the short-term mean itself follows a random walk and is attracted to a long-term mean which remains fixed.

Definitions of the state variables:

r_t = current short-term risk-free rate at time t

m_t = current short-term mean reversion point for the risk-free rate at time t

Definitions of the parameters:

μ = long-term mean reversion point

σ_1 = volatility of the short-term rate

σ_2 = volatility of the short-term mean reversion point

κ_1 = mean reversion strength of the short-term rate to the short-term mean

κ_2 = mean reversion strength of the short-term mean to the long-term mean

Definition of the stochastic process in discrete time:

$$r_{t+1} = r_t + \kappa_1(m_t - r_t) + \sigma_1 W_1$$

$$m_{t+1} = m_t + \kappa_2(\mu - m_t) + \sigma_2 W_2$$

The random shocks W_1 and W_2 are random draws from a Gaussian (Normal) distribution with mean 0 and variance 1.

There are several variations on this model.

- The two random shocks may or may not be correlated.
- The two stochastic processes (two factors) shown above are in a form analogous to the Vasicek model. Each could be changed to be analogous to a Cox-Ingersoll-Ross, Black-Karasinski, or other formulation, with suitable changes to values of the parameters.

Characteristics of this model type that deserve note:

- A variety of yield curve shapes can be generated using a fixed set of parameter values while changing the values of the state variables. For example, a humped yield curve is generated when the short-term mean is greater than either the short-term rate or the long-term mean.
- In some scenarios the short-term mean will remain far from the long-term mean for an extended period of time. Since the short-term rate reverts to the short-term mean, this tends to create scenarios that may be characterized as “low-for-long” or “high-for-long” in greater proportion than with a model with a single fixed mean reversion point.

Reflecting the market price of risk in a double-mean-reverting model

The market price of risk is reflected by the difference in drift due to the difference between real-world and risk-neutral parameter values. The simplest approach is to simply increase the long-term mean reversion point so that the risk-neutral value of the parameter μ is $\hat{\mu} = \mu + \lambda_\mu$. Alternately, one can use the approach discussed earlier where one starts with the formulaic yield curve based on real-world parameters and add a vector of term premiums to reflect the market price of risk. There are, of course, other approaches as well.

THE THREE-FACTOR COX-INGERSOLL-ROSS MODEL

A three-factor Cox-Ingersoll-Ross model is simply the sum of three independent one-factor Cox-Ingersoll-Ross models. The three factors each have their own parameters for mean, volatility, and mean reversion strength. The short-term rate is the sum of the state variables of the three factors.

Definitions of the state variables:

$r_{1,t}$ = current short-term risk-free rate in factor 1

$r_{2,t}$ = current short-term risk-free rate in factor 2

$r_{3,t}$ = current short-term risk-free rate in factor 3

Definitions of the parameters:

Each of the three factors has its own set of the three parameters in a Cox-Ingersoll-Ross model:

κ = mean reversion strength or speed

θ = mean reversion point

σ = volatility

The values of these parameters are different between the factors, so there are nine parameters in total before reflecting the market price of risk.

Definition of the stochastic process in discrete time:

Each of the factors follows the basic Cox-Ingersoll-Ross process using its own state variable and its own set of parameter values. Here is the process used for each factor expressed in continuous time:

$$dr = \kappa(\theta - r)dt + \sigma\sqrt{r}dW$$

Since the short-term rate is the sum of the state variables, the price of a pure discount bond is

$$P_T = E \left[\exp \left(- \int_0^T (r_{1,t} + r_{2,t} + r_{3,t}) dt \right) \right]$$

The mathematical form of the Cox-Ingersoll-Ross model allows derivation of a closed-form expression for this price. That is a significant advantage of this type of model.

Characteristics of this model type that deserve note:

- The three factors in this model are loosely analogous to the three components of the yield curve that arise from principal component analysis. A calibration of this model generally results in the three factors having very different parameter values for the mean reversion strength or speed. The parameter values are low, medium, and high. These correspond to the principal components as follows:
 - Low mean reversion speed corresponds to the “level” component
 - Medium mean reversion speed corresponds to the “slope” component

- High mean reversion speed corresponds to the “curvature” component
- The three factors each have a fixed mean reversion point, so in sum they have a fixed mean reversion point. This makes this kind of model less likely to produce scenarios that remain far from the mean for an extended period of time as in “low for long” or “high for long”.
- The movements of the state variables in this model are assumed to be independent, and scenarios are generated using independent random shocks to each factor. However, when the state variables are fit to historical yield curves, the historical movements of the state variables are significantly correlated. This indicates that while this model can generate reasonable distributions of future interest rates, the path-wise behavior of scenarios may be less realistic than some other models.

Reflecting the market price of risk in a three-factor Cox-Ingersoll-Ross model

As was discussed earlier for the one-factor Cox-Ingersoll-Ross model, it is common to reflect the market price of risk by simply defining adjustments to the parameters θ and κ in each of the factors. That makes six different adjustments (2 per factor x 3 factors). Combined with the 9 parameters already in the model, that makes 15 total parameters in a combined real-world / risk-neutral model¹¹.

As with any mean-reverting model, adding a market price of risk that produces positive term premiums in the generated yield curves is most directly done by an adjustment (increase) in the mean reversion point. In this model each factor has its own mean reversion point and its own mean reversion speed. One can manipulate the shape of the curve of term premiums by deciding how much of the adjustment to the total mean reversion point is allocated to each of the three factors. Putting most of the adjustment in the faster-reverting factors will lead to a curve of term premiums that is steeper on the short end and levels off at the long end, which some actuaries believe is most realistic.

¹¹ That’s before any shift that may be added to allow the model to simulate negative interest rates. The shift would be a 16th parameter.

Appendix C: The meaning of arbitrage-free

Arbitrage is the generation of profit with no risk. It occurs when the same item can be purchased somewhere at one price and sold elsewhere at a higher price. Markets are assumed to be arbitrage-free, so any market consistent or market-coherent ESG must be arbitrage-free.

In a world where countless kinds of derivative securities exist, the detection of arbitrage in any scenario generator is a complex task. While arbitrage is defined as buying and selling the same “item” at different prices, the “item” could be composed of two or more other items, often including a derivative of the item itself. A test for arbitrage must evaluate all equivalent combinations of items. Complicated mathematical tests have been devised for that purpose. Section 9 of *Economic Scenario Generators – A Practical Guide* (Pedersen et. al. 2016) describes such tests.

A simple concept underlies those complicated tests. The concept is that the generated paths of future returns (i.e., interest rates, stock prices) must be consistent with market expectations on the starting date of the scenarios. One of those expectations is the path of the short-term risk-free rate. To be arbitrage-free, an ESG for investments that involve risk must produce some scenario returns higher and some lower than the short-term risk-free rate. That way there are some scenarios that result in gains and some that result in losses, relative to keeping money risk-free in a money market fund. If all scenarios result in gains, or all result in losses, then arbitrage exists, and the generator is not arbitrage-free. To be arbitrage-free, any possibility of gain (earnings higher than the risk-free rate) must be accompanied by a possibility of loss.

This does not mean that the average return must equal the risk-free rate. The average scenario return for a class of investments can be much different from the risk-free rate while still allowing the possibility of both gains and losses and thereby remaining arbitrage-free. This is one of the keys to understanding the connection between real-world and risk-neutral ESGs. In real-world ESGs the average scenario return for an investment depends on the level of risk in the investment. In risk-neutral ESGs the average scenario return for all investments is the same short-term risk-free rate, but that risk-free rate follows an unrealistic average path. Both kinds of ESGs can be arbitrage-free; the difference is that the market price of risk is nonzero in a real-world generator and that leads to average returns on risky investments that differ from the risk-free rate.

Appendix D: Term premiums and the market price of risk

Term premiums are the reward investors demand for taking the risk of locking in a fixed interest rate in a world where interest rates change over time. The risk is that interest rates may rise, so a locked-in rate has both an element of opportunity cost and an element of risk that the market value of a fixed-rate investment may decline. Those two risks are flip sides of the same thing – they are 100% correlated.

Since term premiums are a reward for taking risk, they are a function of the market price of risk. If you can quantify the market price of risk, you should be able to calculate the corresponding term premiums. This appendix explains how to do that.

The market price of risk has been defined as this ratio, where r is the short-term risk-free rate:

$$\text{Market price of risk} = \lambda = \frac{(\text{Expected average return}) - r}{\text{standard deviation of price}}$$

For default-free fixed-income investments, the numerator of the ratio is the term premium. The denominator is the volatility of the price. Recall that in an arbitrage-free model, if the short-term risk-free rate at time t is r_t then the spot price P for maturity $t=T$ is the expectation:

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

The price is a function of the risk-free rate at time 0, which is r_0 . One should be able to estimate the sensitivity of the price to a change in r_0 , which would be noted as $\frac{\partial P_T}{\partial r_0}$. We also can estimate the volatility of r_0 which can be notated as σ_r . We can then quantify term premiums using the following algebra.

$$\lambda = \frac{(\text{Expected average return}) - r}{\text{standard deviation of price}}$$

$$\lambda = \frac{(\text{Term premium})}{\sigma_r \left(\frac{\partial P_T}{\partial r_0} \right)}$$


$$\lambda \sigma_r \left(\frac{\partial P_T}{\partial r_0} \right) = (\text{Term premium})$$

This gives an expression for the term premium given three things:

- An estimate of the market price of risk λ
- The volatility of the short-term interest rate σ_r
- The sensitivity of the price to the short-term rate $\frac{\partial P_T}{\partial r_0}$

Consider the situation where the values of λ and σ_r are constants¹². The sensitivity of the price is a function of the term to maturity and is greater for longer maturities. This means that term premiums can be expected to increase with term to maturity and follow a curve that mirrors the shape of $\frac{\partial P_T}{\partial r_0}$ which varies by term to maturity.

The Cox-Ingersoll-Ross model is one such model and it has a formulaic expression for the spot prices P_T . That makes it possible to calculate the value of $\frac{\partial P_T}{\partial r_0}$ and quantify the term premiums resulting from a set of chosen parameter values. The chart below illustrates term premiums from the Cox-Ingersoll-Ross model based on one reasonable calibration.



Similar curves arise from other mean-reverting models.

When a curve of term premiums is quantified in this way, one needs to be clear about the measure of the observed yield curve to which term premiums correspond. When term premiums are quantified in this way, they are part of the forward rate curve. If they were quantified as part of the spot curve or the par coupon curve, they would typically be lower and less steep on the short end¹³.

The connection between real-world and risk-neutral interest rate models can be seen clearly in this context. Consider the formula for the price of a pure discount bond:

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

¹² There are of course models where λ and σ_r are not constants. These include multi-factor models or models with stochastic volatility. The mathematics associated with such models gets very complex and is beyond the scope of this document. Nevertheless, the same conceptual framework applies. One can gain much intuitive insight from the simple analysis presented here.

¹³ The reader is assumed to know how to transform a yield curve into its three equivalent forms – the par coupon curve, the spot curve, and the forward rate curve. One source is Fabozzi 2021, pages 692-695, 709-713.

The price is defined based on the path of r_t , the instantaneous forward discount rate. We can break the path of r_t into two parts: the term premium p_t and the remainder s_t . The remainder conceptually corresponds to the path of the risk-free short-term rate, which does not include term premiums. We then have:

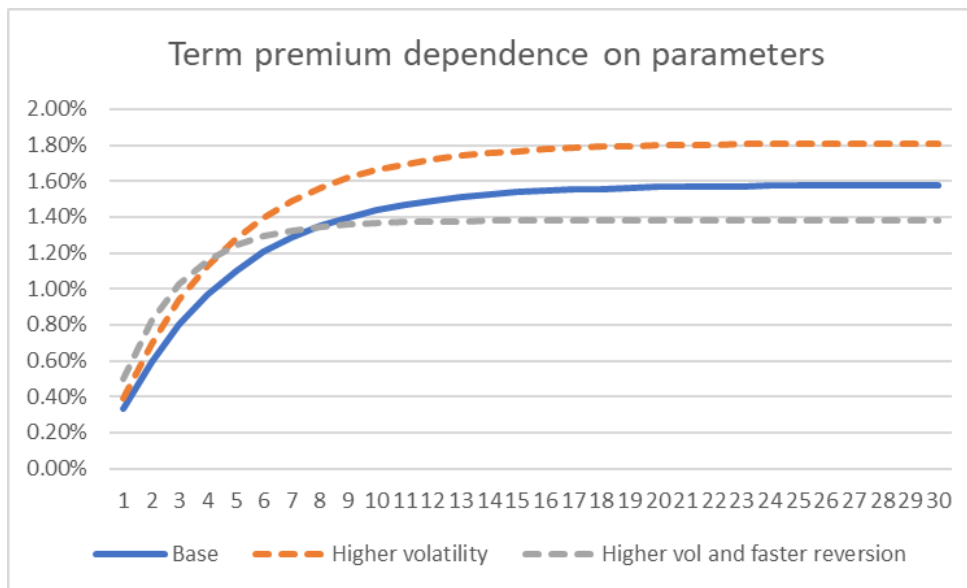
$$P_T = E \left[\exp \left(- \int_0^T (s_t + p_t) dt \right) \right]$$

The risk-neutral parameters for an interest rate scenario generator define the stochastic path for r_t . The real-world parameters for an interest rate scenario generator define the stochastic path for s_t , which is the path of r_t with term premiums removed. The difference between those paths is the path of term premiums p_t .

While term premiums are not directly observable at a point in time, cross-sectional studies over time can provide estimates of term premiums. Such studies have been done, and they suggest that term premiums follow a pattern by maturity that is much like that shown in the chart above, rising quickly for short maturities and leveling off for longer maturities.

NOTES ON CALIBRATION OF TERM PREMIUMS

The chart shown above is based on one reasonable calibration of a model. Term premiums are sensitive to the calibration. The chart below illustrates term premiums based on some alternate parameter values for volatility and mean reversion speed, all using the same market price of risk within the same model.



In more complex interest rate models, term premiums can also be sensitive to the level and shape of the starting yield curve.

One may wish to adapt an ESG so that it can reflect a desired set of term premiums. The exhibit below outlines the steps used to generate scenarios in a market-coherent real-world ESG. The basic process is shown on the left, and the changes needed to reflect pre-specified term premiums are shown on the right.

Step	Basic process	Revised process (specified term premiums)
<i>Initialize state variables</i>	<i>Fit state variables to the spot prices in the observed starting curve using the risk-neutral parameters.</i>	<i>Subtract the specified term premiums from the observed starting curve and fit the state variables to the resulting curve using real-world parameters.</i>
<i>Generate paths of state variables</i>	<i>Use the real-world parameters to generate the stochastic paths.</i>	<i>Same.</i>
<i>Calculate yield curves at each time step based on the state variables at that time</i>	<i>Use the risk-neutral parameters to calculate the yield curve based on the state variables.</i>	<i>Use the real-world parameters to calculate a yield curve based on the state variables, and then add the specified term premiums to that curve.</i>

Note that there is no use of risk-neutral parameters in the revised process. The use of pre-defined term premiums replaces the role of the risk-neutral parameters. This emphasizes the fact that the role of risk-neutral parameters in a real-world market-coherent ESG is simply to define the term premiums. If term premiums are provided explicitly in another way, then there is no role for risk-neutral parameters in generation of real-world market-coherent scenarios.

The model under discussion here assumes that the yield curve is governed by two forces – the future expected path of the short-term rate and the term premiums based on the market price of risk. The observed yield curve can also be influenced by other things such as supply and demand factors and inflation expectations. These issues need to be considered when evaluating the fit of a model to historical data.

Appendix E: Fair value, Principle-based Reserves, and the Valuation Manual

The valuation of insurance reserve liabilities (“reserves”) for accounting purposes has evolved in the 21st century so that it is often based on stochastic simulation and discounting of future cash flows. In the U.S. such valuation methods have been called “principle-based” (Center for Insurance Policy Research 2021) and are encoded in the National Association of Insurance Commissioner’s (NAIC) Valuation Manual for use in required regulatory financial reporting. Internationally, a different and new standard of accounting for insurance contracts has also employed such methods (IFRS 2021), but with important differences.

There is a common principle on which all such methods are based. The principle is that the value of a block of insurance contract reserves is equal to the value of the assets required on the valuation date to pay them off, with some degree of uncertainty or margin for risk. This principle ties the valuation of the insurance contract reserves to the valuation of the assets. The largest difference between reserve valuations in the international accounting standard and the U.S. statutory Valuation Manual is that the valuation of assets is at market in the international standard while it is largely based on amortized cost in U.S. insurance regulatory (“statutory”) reporting.

There are two ways to tie the value of the reserves to the value of the assets in this kind of principle-based approach. One way is to directly determine the amount of assets needed on the valuation date to pay off the contracts and set the value of the reserve equal to the value of those assets. The other way is to determine the rate of investment return in each time step of an asset-liability cash flow simulation and discount the liability cash flows at a rate equal to the path of projected investment returns. If the amount of starting assets in the simulation is just sufficient to run out when the last contract payment is made, then these two methods should produce identical results¹⁴.

Different accounting bases prescribe different asset valuation methods, and different asset valuation leads to the reporting of different investment yields. When the reserve is calculated by discounting cash flows at the projected investment return on the assets, one must use the investment return as constrained by the applicable accounting standard. The reported investment return may differ from the return on market value in cases where assets and reserves are not held at a market value. For example, the investment return as measured under U.S. statutory accounting can be substantially different from current returns based on market value.

Stochastic valuation of reserves with an asset-liability simulation model involves calculation of a separate liability value for each stochastic scenario and then using the set of scenario-specific values to determine a single reported value. The next two subsections discuss methods commonly used to get scenario-specific values and to determine the single reported value.

OBTAINING SCENARIO-SPECIFIC VALUES

When using stochastic scenarios for valuation of reserves in an asset-liability model, the amount of starting assets required to fulfill the contract differs for each scenario, unless the asset portfolio is a true replicating portfolio. True replicating portfolios are rare, so one must deal with the question of how to set the amount of starting assets.

As a practical matter, one does not use a different amount of starting assets for each scenario. This creates two issues. First, one must determine what amount of assets to start with, and second, one must have a method to

¹⁴ Significant differences in results from these two approaches can arise when the path of investment returns has not been determined correctly based on the accounting standard in use.

calculate the value of the reserve when the assets do not exactly satisfy the contract cash flows in a particular scenario.

A generally accepted approach is to start with an amount of assets very close in value to the ending single reported value of the reserve. This approach is somewhat circular and may require a bit of trial and error to arrive at that value. Nevertheless, this approach is required in the Valuation Manual, where the amount of starting assets is required to be within a specified narrow range around the reported reserve amount. Generally, this means that the starting assets will be more than sufficient to fund the reserve in most scenarios, but there may still be some particularly adverse scenarios where the assets run out before all contract obligations are paid. In those scenarios, the asset-liability model simulates borrowing to pay the obligations, and the investment return becomes the rate of interest simulated as paid on the borrowed funds.

If the same amount of assets is used at the start of each scenario, then the scenario-specific value of the reserve cannot be set equal to the value of the starting assets. Therefore, the approach of discounting the cash flows at the path of investment returns (as modeled scenario by scenario and as reported under the applicable accounting standard) can be used to determine the scenario-specific value of the liability.

In the Valuation Manual, another method is sometimes required; it is the “greatest present value of accumulated deficit” (GPVAD). Under this method, one determines a deficit defined as the excess of the outstanding liability over accumulated assets at each time step in the projection¹⁵. Each of those amounts is discounted back to the valuation date, and the GPVAD is the most negative of those discounted values. The scenario-specific reserve value is then defined as the sum of the value of the starting assets and the GPVAD. In a scenario where the starting assets are exactly sufficient to pay off the obligations, the GPVAD is zero and the scenario-specific reserve value is equal to the value of the starting assets.

The discount rate used for the GPVAD theoretically reflects the expected return on the extra starting assets that would be required to fully fund the obligation and eliminate any deficits in a particular scenario. For regulatory purposes that could be the same as the projected return on the assets in the asset-liability projection, or it could be set conservatively at a level close to the risk-free rate. In either case, it is important to understand that this discount rate is not used to discount cash flows; it is only used to discount accumulated deficits. The biggest part of the liability value is set equal to the value of the starting assets, with the GPVAD often being a relatively small adjustment that varies by scenario.

BLENDING SCENARIO-SPECIFIC VALUES INTO A SINGLE REPORTED VALUE

When the accounting basis for asset valuation is market value, then risk-neutral methods can be used, and the single liability value is the simple average of the scenario-specific values.

Any valuation of insurance contracts needs to include a margin for risk. Risk-neutral methods include a margin implicitly because risk-neutral calibration gives added probability weight to adverse scenarios. Whenever real-world scenarios are used, the simple average of the scenario-specific values does not include such an implicit margin and a margin must be added in some other way.

The two most common methods for adding such a margin in the context of real-world scenarios are these:

¹⁵ Since it can be difficult or time-consuming to determine the outstanding liability value at each time step in a scenario, an approximation or “working reserve” may be defined for this purpose. Sometimes the “working reserve” is simply zero and the deficit is simply the negative of the amount of accumulated assets.

- **Contingent tail expectation (CTE).** Under this method the scenario-specific values are sorted from least to greatest, and only the greatest are included when calculating an average value. The number of scenarios included in the average is a specified percentage of the total. The Valuation Manual specifies the “70 CTE” which means that the smallest 70% of the scenarios are excluded from the average and only the largest 30% are included.
- **Cost of capital method.** Under this method the scenario specific values are each increased by adding an imputed cash flow equal to the cost of capital in each time step of the scenario. The margin is the present value of those imputed cash flows. The cost of capital is intended to represent the market price of risk. In theory this approach may be considered consistent with market pricing but estimating the cost of capital is felt by some actuaries to be problematic.

SELECTING A SCENARIO GENERATOR THAT IS FIT FOR PURPOSE

If the accounting basis for assets is market value, then the scenarios used for liability valuation should be able to reproduce market prices of traded assets. While risk-neutral scenarios are always calibrated to do that, real-world scenarios may or may not be calibrated that way.

When using a real-world ESG it is important that the calibration approach be appropriate for the accounting basis. Real-world scenarios that are not calibrated to reproduce market prices may be appropriate when the accounting basis is not market value.

Recall that valuation using real-world scenarios must include an add-on margin for risk. When the scenarios used are not market-coherent, then the margin may arise partly from conservatism in the scenario calibration and partly from an explicit margin. That is the case with the Valuation Manual. U.S. regulators mandate the use of a real-world scenario generator that is not market-coherent and then add a margin using the CTE approach.

A special issue arises when hedging is to be simulated within each scenario in the asset-liability model. Hedging depends on market-consistent¹⁶ valuation, including “Greeks”. To simulate hedging activity in each time step one needs market-consistent values for both the hedges and the contractual obligations being hedged. Such market-consistent values can be obtained through separate stochastic valuations at each time step. Stochastic valuations at each time step in a long-term simulation are often called “inner loop” valuations to indicate that the scenarios used for valuation are separate and different from the “outer loop” scenario that defines the conditions at each time step. The starting date for a set of “inner loop” scenarios is a valuation date anywhere within time span of the outer loop scenario. Since inner loop valuations are often used to determine a market-consistent value for derivatives used for hedging, they often employ risk-neutral scenarios calibrated on the fly to the yield curve at their starting date. This is true even if the outer loop scenario is real-world and neither market-consistent nor market-coherent in nature. In situations involving hedging, a risk-neutral ESG may be used for the inner loop valuations even though the outer loop is based on scenarios from a real-world ESG. That can be appropriate, for example, when simulating a clearly defined hedging strategy in a valuation that complies with the Valuation Manual where the outer loop scenarios are real-world and neither market-consistent nor market-coherent.¹⁷

¹⁶ In this paragraph the term “market-consistent” means a valuation technique that reproduces observed market prices. Both “market-consistent” and “market-coherent” valuations satisfy this definition.

¹⁷ See Question 12.1 (Slutsker 2019).

References

Center for Insurance Policy and Research. Principle-Based Reserving (PBR). *National Association of Insurance Commissioners (NAIC.org)*. Last updated May 28, 2021. <https://content.naic.org/cipr-topics/principle-based-reserving-pbr>

Fabozzi, Frank. 2021. *Fixed Income Securities*. Ninth Edition. McGraw Hill.

IFRS. 2021. *IFRS 17 Insurance Contracts Standard 2022 Issued*. International Financial Reporting Standards Foundation. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-17-insurance-contracts/>

Pedersen, Hal, Mary Pat Campbell, Stephen L Christiansen, Samuel H. Cox, Daniel Finn, Ken Griffin, Nigel Hooker, Matthew Lightwood, Stephen M Sonlin and Chris Suchar. 2016. *Economic Scenarios – A Practical Guide*. Chicago: Society of Actuaries. <https://www.soa.org/globalassets/assets/Files/Research/Projects/research-2016-economic-scenario-generators.pdf>

Slutsker, Benjamin and Dylan Strother, Cindy McGovern et al. 2019. *Principle-Based Approach Projections Practice Note*. Washington. American Academy of Actuaries. https://www.actuary.org/sites/default/files/2019-12/PBA_Projections_Practice_Note.pdf

About The Society of Actuaries Research Institute

Serving as the research arm of the Society of Actuaries (SOA), the SOA Research Institute provides objective, data-driven research bringing together tried and true practices and future-focused approaches to address societal challenges and your business needs. The Institute provides trusted knowledge, extensive experience and new technologies to help effectively identify, predict and manage risks.

Representing the thousands of actuaries who help conduct critical research, the SOA Research Institute provides clarity and solutions on risks and societal challenges. The Institute connects actuaries, academics, employers, the insurance industry, regulators, research partners, foundations and research institutions, sponsors and non-governmental organizations, building an effective network which provides support, knowledge and expertise regarding the management of risk to benefit the industry and the public.

Managed by experienced actuaries and research experts from a broad range of industries, the SOA Research Institute creates, funds, develops and distributes research to elevate actuaries as leaders in measuring and managing risk. These efforts include studies, essay collections, webcasts, research papers, survey reports, and original research on topics impacting society.

Harnessing its peer-reviewed research, leading-edge technologies, new data tools and innovative practices, the Institute seeks to understand the underlying causes of risk and the possible outcomes. The Institute develops objective research spanning a variety of topics with its strategic research programs: aging and retirement; actuarial innovation and technology; mortality and longevity; diversity, equity and inclusion; health care cost trends; and catastrophe and climate risk. The Institute has a large volume of topical research available, including an expanding collection of international and market-specific research, experience studies, models and timely research.

Society of Actuaries Research Institute

475 N. Martingale Road, Suite 600

Schaumburg, Illinois 60173

www.SOA.org