CAS RESEARCH PAPERS

A USER'S GUIDE TO ECONOMIC SCENARIO GENERATION IN PROPERTY/CASUALTY INSURANCE

Sponsored by
Casualty Actuarial Society
and Conning

A User's Guide to Economic Scenario Generation in Property/Casualty Insurance

Conning

Citation:

Conning, "A User's Guide to Economic Scenario Generation in Property/Casualty Insurance." *Casualty Actuarial Society*, CAS Research Papers, 14 Oct. 2020, www.casact.org/research/research-papers/Users-Guide-Economic-Scenario-Generation-in-P&C-Insurance.pdf.

Contents

A User's Guide to Economic Scenario Generation in Property/Casualty Insurance

Conning

Introduction

An economic scenario generator (ESG) is a computer-based model that provides many simulated examples of possible future values of various economic and financial variables. Those scenarios, along with analysis of the stochastic distribution of scenario outcomes, illuminate the nature of risk elements within the economy that drive financial variability. As such, an ESG can provide insights into the relative advantages and disadvantages of alternative operating and strategic decisions.

An ESG typically comprises several interacting modules, although the specific nature of the modules and their interrelationships may vary from one ESG to another. Like any model, an ESG model can be characterized by three major parts: input, output, and the calculations that go on in between. Typically, one variable tends to serve as a driver of the other variables being generated in a scenario. Compared with deterministic economic scenarios, econometric models, and macrofinance models, an ESG simulation can deliver a better view of scenario probabilities, a broader range of scenario outcomes, and greater complexity of scenarios. Modeling can follow either risk-neutral or real-world approaches: risk-neutral (or market-consistent) frameworks are required by certain regulatory authorities for valuation of insurance liabilities, while real-world modeling is appropriate when projecting future values of economic and financial variables.

We intend this publication to serve as a basic guide to ESGs, with an emphasis on applications for the property/casualty insurance industry. The first half of the publication provides general information on the nature and applications of ESGs and discusses their specific applications in the insurance industry. It also discusses essential features of a good ESG and offers guidance on stochastic processes and modeling of certain economic and financial variables. We discuss financial market model specification, model calibration, and model validation and their importance in ensuring that the ESG will render simulation results that are relevant and sufficiently robust and that realistically reflect market dynamics.

In the second half of the publication we illustrate how one group of researchers approached the development of an ESG, describing issues and decisions made in constructing and using that specific ESG. The second half also discusses sources of data and illustrates a validation process using the model to visualize outcomes and support recalibration. Specific considerations relating to the projection time frame (short horizons versus longer horizons) are explored in depth—these are particularly relevant in the calibration process of ESGs in the property/casualty environment.

Finally, we discuss the range of choices a user has for ESG development software, contrasting open-source ESGs with solutions available from commercial vendors. The publication closes with an annotated bibliography of literature in the field as a guide for further research.

Executive Summary

Chapter 1: What Is an Economic Scenario Generator?

An economic scenario generator (ESG) is a computer-based model that provides many simulated examples of possible future values of various economic and financial variables. Those scenarios, along with analysis of the stochastic distribution of scenario outcomes, illuminate the nature of risk elements within the economy that drive financial variability. As such, an ESG can provide insights into the relative advantages and disadvantages of alternative operating and strategic decisions.

An ESG is typically used in combination with other models—components that use the economic scenarios as inputs and then calculate items of organizational interest. An ESG's value is in its ability to simulate and project economic scenarios in a structured and rigorous way. This is necessary because financial and economic variables are stochastic—they change over time in a largely unpredictable way.

An ESG is typically composed of several interacting modules, although the specific nature of the modules and their interrelationships may vary from one ESG to another. Like any model, an ESG model can be characterized by three major parts: input, output, and the calculations that go on in between. An ESG typically simulates all relevant economic and financial variables, but one variable tends to serve as a driver of the other variables being generated in a scenario. Projections of variables are developed on a holistic basis, and asset classes are covered. Finally, parameter values can be updated by the user, and the ESG can be assessed and validated.

Two critical aspects of modeling the financial and economic variables in the modules of the ESG include the parameterization and calibration of the variables and the inclusion of proper relationships of correlation and other interrelationships between the variables. The interrelationships may be developed by a cascade structure, correlation mechanisms, and direct linkages.

For ESG applications, modeling can follow either risk-neutral or real-world approaches. Some regulatory authorities require risk-neutral (or market-consistent) frameworks for valuation of insurance liabilities. Real-world modeling is appropriate when projecting future values of economic and financial variables. Modeling can also follow discrete-time and continuous-time mathematics. Generally, continuous-time modeling leads to more convenient mathematics.

Analytical tasks may also distinguish between arbitrage-free and equilibrium models. Arbitrage-free requires that the relationships between economic and financial values do not allow for the possibility of arbitrage. An equilibrium model specifies that the interest process balances supply and demand, which is often better for longer-term horizons.

Compared with deterministic economic scenarios, econometric models, and macrofinance models, an ESG simulation can provide a better view of scenario probabilities, a broader range of scenario outcomes, and greater complexity of scenarios. However, an ESG can appear as a black box, may contain significant model risk, and can require significant resources for maintenance and development. With increased computer resources and better and more detailed data, the greater scope and sophistication of ESGs has resulted in more resources devoted to data and to the general maintenance and operation of the ESG. Thus, it is critical that the design of an ESG carefully consider the objective of the analysis to be undertaken, with sufficient attention being paid to transparency and documentation.

Chapter 2: Applications of Economic Scenario Generators

ESGs are a critical component of a wide range of applications used by insurers in managing the economic risks of their operations. For a given application, it is critical that the ESG be suitable and properly maintained relative to the application's purposes.

The most common ESG-driven applications for property and casualty are asset–liability management (ALM) systems (used in assessing, establishing, and monitoring investment strategies) and economic capital systems (used to calculate and monitor economic capital).

ALM systems deal primarily with economic risk mitigation, in which the range of adverse economic events is narrowed or reduced while still maintaining a healthy likelihood of positive investment growth.

Economic capital systems typically focus on shorter time horizons and involve significantly more scenarios in order to establish reliable tail metrics.

Any ESG application will have some practical limitations based on its underlying ESG and any functionality (trading strategies, etc.) that has been provided/implemented in support of the application use case(s). Users of an ESG application must appreciate any such limits in order to appreciate how best to interpret and communicate results generated.

Chapter 3: Nature and Role of ESGs in Property/Casualty Insurance

For property/casualty insurers, the ability to assess financial statement values, as well as the impact of operational or strategic decisions, requires being able to enumerate and describe a wide range of the possible states of economic and financial conditions. Some of the more important variables that a P/C insurer should consider when building an ESG include the valuation of assets and liabilities, economic capital and regulatory requirements, strategic and operational decision-making, and risk management.

Investment portfolio decisions may be based on regulatory requirements as well as the need for maintaining a certain level of liquidity. General characteristics of P/C insurers, including prospective cash flows in the context of a going-concern enterprise, can dictate many of their asset and liability cash flow patterns, and consequently their asset–liability management decisions. Asset and liability portfolio values may be influenced by financial factors such as interest rates (risk free, risk premia, and term premia), credit risk (credit rating migration, default risk intensity), inflation (general and line-of-business specific), equity returns, and mortgage delinquency and prepayment patterns. These characteristics, along with the specific attributes and business models of individual companies, and the purpose for which the model is designed, dictate the kinds of economic and financial variables that should populate an ESG.

There are several points of intersection between P/C underwriting and operational results and the economic and financial variables an ESG generates. For example, premium volumes and losses associated with many P/C lines of business are related to economic conditions, often causally. Furthermore, underwriting and operating factors tend to undergo significant cyclicality from periods of high premium rates and low loss ratios to low premium rates and high loss ratios. Thus, the ability to model a P/C insurer relative to a range of different economic conditions over time is critical.

Valuation of the reserves for outstanding losses (the largest liabilities of a P/C insurer) is largely the purview of actuaries. While the reserve shown on the insurer's balance sheet is a single "best-estimate" value, the loss reserve is actually a stochastic value with variability around the best estimate, and the best estimate may itself vary under different scenarios or conditions. A good ESG provides an actuary with a robust tool to build deeper insight into the potential volatility of future loss payments.

Other important factors in P/C balance sheet considerations include the volatility of assets (and the leverage of invested assets against surplus), the impact of foreign exchange models and multi-economy factors, and the effect of different time horizons on different line-of-business models with variable claims payout periods.

Some aspects of asset risk can be evaluated through a strategic asset allocation analysis. An important aspect of strategic asset allocation is developing an efficient frontier of investment classes to optimize risk and return. For example, assessing the duration behavior of the investment portfolio against the duration of liabilities on the balance sheet throughout a range of economic scenarios can lead to a deeper understanding of the effect of interest rates and other economic factors on assets, liabilities, and surplus.

Economic capital and regulatory requirements for P/C insurers tend to be influenced by extreme tail events, requiring responses in the form of stress testing. Often, extreme events can influence multiple aspects of the business—such as, for example, catastrophic events that influence the general health of the economy—leading to a potential double impact on the P/C insurer. Inflation could also accelerate due to supply-and-demand issues after a major catastrophe. This is precisely the type of application at which a good ESG can excel.

Analysis of extreme events can also influence strategic and operational decision-making. An ESG cannot itself make decisions about strategic or operational alternatives, but it can provide a consistent basis for evaluating the impact of a decision across a range of different possible future circumstances.

Application of a consistent model can also provide insight into the cost and risk trade-offs of risk management questions and potential solutions.

Done correctly, an ESG can provide foundational information for making many types of corporate decisions, but to take full advantage of this modeling information, it is critical that, across an entire corporate model—for example, an enterprise risk management model—the various modules making up the full model be consistent with one another.

Chapter 4: Perspectives on Developing and Maintaining an ESG

An ESG is a collection of models under a defined architectural structure that incorporates a specified level of detail and a selection of appropriate stochastic dynamics. Development and maintenance of an ESG require a careful approach to parameterization and a disciplined maintenance process to adequately reflect both historical and prospective financial dynamics.

ESGs are developed under a specific and intentional architecture that accommodates the appropriate interaction of component models. The way in which the component models are structured and interact with one another affects the causality and correlation structure of the ESG as well as the calibration methodology. This often employs a cascade structure and vector autoregressions. Interest rates and total returns are generally key output components.

When ESG simulation output is generated, an assumption needs to be made as to the level of detail to be included, including simulation frequency. ESG variables such as interest rates need to be computed and stored with a tenor structure. Price/income relationships, cash flow structures, prepayment features, and default events may also need to be incorporated. As a general rule of thumb, the greater the level of detail in an ESG, the slower it will run, and the larger the simulation data set that is stored.

Development is guided by the stylized facts and institutional details of the key economic variables to be modeled, as well as the level of detail needed for the application. Examples of the stochastic dynamics of a three-month Treasury bill and S&P daily return characteristics are illustrated, including pathwise characteristics such as jumps and volatility.

Data sources for economic modeling can vary in cost and availability. Examples of sources include Bloomberg, Thomson Reuters, Global Financial, Barclays Capital Live, central banks, and bond-rating services, among others.

Parameterization/calibration is a process of selecting model parameters based on certain criteria. Estimating model parameters with historical data and calibrating models to specific market conditions are key parts of the process. The parameterization/calibration process and methodology often reflect a view, both explicit and implicit, with choices dependent on the application—business, regulatory, stress testing, investment management, and so forth.

The validation process involves checking that the calibrated model performs in line with the calibration criteria and that the general behavior of the model is consistent with the stylized facts.

Checking that the model is performing in line with the calibration criteria will usually include a comparison of the simulated model statistics with the calibration targets. Back-testing can provide useful insight into the robustness of models.

The process of ongoing maintenance of an ESG is based on the way in which the models, calibrations, and validation processes interact over time. Under typical conditions, one updates initial conditions (i.e., market data) for each new simulation period, but one does not re-parameterize all the models of an ESG every period.

Chapter 5: What Makes a Good ESG?

An ESG is a complex system and one that must evolve in response to changes in market fundamentals and regulatory requirements. A good ESG has some general characteristics that include the following:

- A good ESG has a sound foundation for the way the models are built and the way the variables are interrelated. It has a full range of modeled financial variables and multi-economy capability.
- A good ESG is capable of accommodating many types of calibration views across a wide variety of benchmarks. A good ESG produces simulation results that reflect a relevant view—i.e., one that is consistent with historical facts.
- A good ESG produces some extreme but plausible outcomes, which encapsulate historical behavior but do not stray too far from market norms.
- A good ESG embeds realistic market dynamics. This requires agreement on a selection of stylized facts and institutional details.
- A good ESG is computationally efficient and numerically stable.
- A good ESG can meet the requirements of regulators and auditing firms.
- A good ESG has fast and robust recalibration capabilities.

Statistical criteria are also important in assessing the quality of an ESG. Statistical calibration criteria are usually numerically specified but can also be qualitative in nature. Statistical criteria belong to one of two broad categories:

- *Qualitative features*. An important first step in validation is to check that the most important qualitative stylized facts are satisfied by the simulated output.
- *Quantitative measures*. Tabular calibration targets will usually include targets for average levels and volatilities.

A path represents one possible future evolution of the economy and therefore represents one possible complete future "economic experience." The importance of pathwise model behavior is that it is the simulated path that represents the way an insurance company will experience the evolution of the economy. If the overall distribution of returns for an asset class is correct but the pathwise behavior does not correspond to the nature of the fluctuations that we see in the historical record, then the model has an issue.

The fundamental process for real-world validation involves comparing calibration criteria against simulated model performance. The criteria used are both qualitative and quantitative. The chapter provides examples of both kinds of criteria applied to several kinds of situations.

Chapter 6: Stochastic Processes and Dynamics for ESG Modeling

In this chapter, we explore ways of modeling certain economic and financial variables. For illustration, we make reference to the financial scenario generator created by Ahlgrim, D'Arcy, and Gorvett (ADG economic scenario generator).

First, we discuss the basics of stochastic processes, from simple discrete random variables to continuous-time processes. The mathematics in this section provides the foundation for the illustrative models introduced in Section 6.3.

The discrete-time framework involves values of variables only at certain points in time, but there is often value in describing economic and financial variables as continuous-time processes. A continuous-time framework can describe a variable's underlying dynamics, but a discrete-time analogue of that continuous-time specification is actually employed when parameterizing and calibrating the model. Finally, one other process that can be useful in modeling economic and financial stochastic processes is a jump process.

Several econometric and statistical techniques can be of value for at least two purposes: (1) modeling the relationships between variables and (2) parameterizing and calibrating the ESG models. One or more key variables may be modeled with all other variables modeled effectively by linking to, or cascading from, those key variables. Some possibilities for achieving such simulated variable relationships include direct linkages, correlation of volatility terms, correlation matrices, and time-varying correlations. In addition, a good model should help the user achieve a better understanding of the interrelationships between variables. These tasks can be supported by a variety of econometric and statistical techniques and considerations, such as, for example, maximum likelihood estimation, Kalman filtering, leading and lagging correlations, and vector autoregression, which involves the modeling of multivariate time series data.

Certain techniques and model requirements are more appropriate for modeling specific economic variables. When deciding upon a model of the term structure of interest rates, two important considerations are the type of model to use (general equilibrium or arbitrage-free), and whether to use a one-factor model or a multiple-factor model. The number of stochastic factors in a term structure model indicates how many "sources of uncertainty" there are. An important characteristic of historical equity returns is the fact that they exhibit "fat tails"—that is, the probability of extreme values well above and well below the mean is greater than would be implied by assuming that the returns follow a normal (or Gaussian) distribution. Inflation is one of the most visible and closely watched of economic variables, and there are numerous ways that it can be and has been modeled—including some approaches that are really quite simple mathematically. For example, because inflation tends to exhibit persistence through time, a simple autoregressive process involving the current value and one or more lagged values has modeling appeal. Other types of regression specifications are also common in modeling inflation. Autoregressive processes may also be useful in modeling other variables, such as equity dividend yields and unemployment.

Finally, we briefly identify a number of important considerations when building and using an ESG, including the sensitivity of data period to stability and responsiveness, a clear understanding of the philosophy or purpose of the ESG, the choice of risk-neutral or real-world framework, the importance of simulation or scenario analysis, and whether negative values will be permitted in modeling certain economic variables.

Chapter 7: Illustrative Modeling of Three Key ESG Components

In this chapter, we discuss modeling issues associated with three key ESG variables: inflation, interest rates, and equity returns.

Although we use the ESG developed by Ahlgrim, D'Arcy, and Gorvett (the ADG model) as a starting point for our discussion of each variable, we go beyond ADG in this chapter to

- address one or two illustrative issues associated with the parameterization and calibration of a model for those variables; and
- identify some alternative modeling structures for those variables.

The ADG model serves as an illustration of how one team of researchers approached the development of an ESG. The model stems from a joint Casualty Actuarial Society/Society of Actuaries request for proposals and has, from the beginning, been intended for an actuarial/insurance audience. It is publicly available, as is its documentation, and is relatively simple compared with some proprietary ESGs. It is built in Microsoft Excel and is run with the simulation software @RISK, an Excel add-on.

We use the ADG model, as well as its structure and development, to illustrate some of the issues and decisions made with respect to its construction and use. Since the ADG model is built in Excel, the logic and algorithmic progression of the model are transparent and can be followed by paging through the various sheets of the workbook.

With a publicly available resource such as the ADG model, it is incumbent upon a user to determine whether the parameterization and calibration underlying the model remain adequate. To illustrate, we look at inflation as an example of re-parameterization.

The ADG model uses a continuous-time version of the stochastic mean-reverting process to simulate the dynamics of inflation. Not surprisingly, the parameter estimates for the inflation process vary greatly depending upon the data period used in the regression. If an ESG is supporting short-run operational decision-making, a shorter and more recent data period may better reflect the dynamics of economic and financial variables for that purpose. If long-run strategic planning is the primary focus of the ESG, then a longer data period including several economic or business cycles may provide a better range of variable values.

The ADG model employs a double-mean-reverting two-factor Vasicek model for the term structure of real interest rates. The two factors represent the short-term real interest rate r and the long-run mean l. Nominal rates are then determined by combining projections of the real interest rates with those for inflation.

The third key variable, equity returns, is modeled in ADG as a regime-switching process. This regime-switching model, based on the approach of Hardy (2001), involves two regimes, high volatility and low volatility, where the likelihood of staying in the current regime or switching to the other regime is governed by transition probabilities, as in a Markov chain. In ADG, the regime-switching equity returns model is used for simulating *excess* equity returns above a risk-free nominal interest rate. Full equity return rates are then simulated by combining the excess return with the simulated risk-free nominal interest rate. Stochastic volatility models with jumps represent a good alternative to regime-switching models.

The validation process of an ESG is described and discussed in Chapters 4 and 5. In this chapter we provide some brief illustrative examples of possible retrospective validation efforts and some general guidance on the verification and validation processes. We examine simulations of inflation variables and look at ex post indications of emerging stock return data to assess the original calibration of that module. We show the value of visualization and the importance of a holistic view encompassing all modules.

We obtained the data we use as the basis for the inflation regressions online from the U.S. Bureau of Labor Statistics. This can be a good source for other economic data such as data on employment and unemployment, pay and benefits, and so on. Other example sources of data are the Federal Reserve Bank of St. Louis Economic Data (FRED) (data on interest rates, inflation, measures of economic activity, real estate, currency, etc.) and Yahoo Finance (downloadable data for individual stocks and stock indices).

Chapter 8: Considerations Related to the Projection Time Frame (Simulation Horizon)

Choosing the specific simulation horizon for an ESG is tightly connected with the overall use to which the ESG scenarios will be applied. Quite short horizons (e.g., one year) can, for example, be seen for economic capital simulations that are more focused on the tail aspects of the distributions. Longer simulation horizons are more common in applications such as asset–

liability management or strategic asset allocation, where applications are focused around the central properties of the distributions, such as mean and standard deviation.

When running an ESG over a long-term horizon the stability of the results is quite important to the user. Ideally, long-term decisions should not be significantly affected by short-term market changes. The best way to achieve this stability is to calibrate the ESG to stable long-term targets. The speed of convergence to the long-term targets can also be important, since that will determine how much influence the starting conditions will have on the overall simulation results. Another important aspect to consider when running ESGs for long-term horizons is whether the downstream system can handle unexpected results—such as negative yields or very high yields—that the ESG produces.

When running ESG scenarios over a short-term horizon a common application is risk-focused analysis. Therefore, the tails of the distributions will be more important when calibrating the ESG for such an application. Yet some of the historical data may refer to time periods with a very different market situation that, therefore, can be of no or very limited use for establishing the calibration targets. And having established the targets for a risk-focused calibration, one still needs to monitor changes in market conditions to ensure that the targets remain appropriate. We have two possible change drivers in an ESG used for short-term risk-focused calibration: (a) change in targets/calibration of the ESG, and (b) change in the starting market conditions.

When one wants to unite both short-term and long-term criteria in one ESG scenario, the ESG should generate percentile-focused results at the start of the simulation and mean and standard deviation–focused results some time into the simulation. However, the distributions that one is focusing on in the short term can be quite different from the ones of interest in the long-term simulations.

We can separate the models used in practice into two broad classes: (a) where the volatility of interest rates is not dependent on the level of interest rates, and (b) where the volatility of interest rates rises with the level of interest rates. For the first class of models, it is impossible to construct an ESG where both short-term and long-term calibration criteria will be met. For the second class of models, there is a risk that the model parameters, when calibrated to the short-term criteria, will either overestimate or underestimate the long-term distribution parameters. However, the possibility at least exists of finding a suitable compromise between short-term and long-term calibration criteria in this class of models.

Chapter 9: Calibrations for One-Year and Short-Horizon Capital Models

There are some circumstances in which an ESG user needs a model designed and calibrated specifically for the projection of assets and liabilities at short simulation horizons. These applications include regulatory capital calculations as well as capital models used for general risk management and investment decision-making.

In this chapter, we introduce the most important questions to consider when implementing an ESG for short-time-horizon simulations. There is no consensus on how a short-term calibration should be defined and implemented. We introduce the most commonly seen approaches and discuss in some detail the relative merits and limitations of each. The reader is advised that the final decision will depend on the application, but also on the models and systems available to him or her. Due to the likely requirement for ongoing maintenance and recalibration of the ESG, the expertise and availability of personnel with particular skill sets may also influence the details of the process chosen.

There are three areas to consider at the outset: the establishment of suitable benchmarks, the task of finding appropriate model parameters that satisfy those benchmarks, and maintenance to ensure that the benchmarks and models remain appropriate.

Establishing short-term calibration benchmarks is arguably more difficult than for the longer time horizon case, because we observe significant variability in the distributions and distributional properties of financial variables over different short time horizons. Two considerations are important when establishing calibration benchmarks: which risk metrics should be the focus of the target-setting procedure, and what methodology should one use to set the numerical value of the benchmark statistic?

Most applications will require calibration benchmarks for all or some of the following:

- *Initial conditions* (often called time-zero values), which act as a starting point for all simulations
- A *best estimate,* which most typically corresponds to the mean or median of the simulated distribution
- A *volatility benchmark* giving a target for the standard deviation of the simulated distribution
- *Relevant extremes* that the simulated distribution should exhibit

The setting of model benchmarks would ideally also take into consideration the features and limitations of the available models. The propensity for markets to be volatile puts greater emphasis on model and system selection for short-term calibrations than it does for the long-term case. When selecting an ESG for the purposes of short-term calibration, practitioners should consider the model's flexibility, sufficiency of calibration tools, and extensibility to augment deterministic scenarios, as well as the level of expertise required to maintain the calibrations.

Chapter 10: Software for Economic Scenario Generation

Users of ESGs have a wide choice of ESG software solutions. Open-source ESGs are available with relatively limited functionality. Commercial vendors offer a wide range of ESG software solutions. Commercial ESG software is widely available in the cloud, which can offer speed and cost-saving advantages.

The choice of an ESG software solution will be driven by the company use case, anticipated needs, and information technology requirements. A solid understanding of the issues involved in the development and maintenance of an ESG and the attributes of a good ESG are important guideposts in selecting the appropriate ESG software.

Chapter 11: Guide to the Literature on Economic Scenario Generation

In vibrant research areas such as those related to ESGs, it is difficult to thoroughly survey and identify the literature, since the pool of relevant and potentially meaningful papers and research reports is ever growing. The best one can do is provide an annotated bibliography of the particularly significant or "classic" papers, a list of additional papers with which researchers may want to become familiar, and a few suggestions for where future literature of relevance may be found. These are our goals in this chapter.

Section 11.1 identifies the characteristics of the ESG literature as a whole, and the contexts within which potentially useful material may be found. Section 11.2 describes several of the major, classic papers that deal directly with ESGs. These papers include material with which every designer, builder, and probably even user of ESGs should be familiar.

Section 11.3 also provides an annotated bibliography, of major papers categorized according to their particular subject matter. Deeper understanding of this material is appropriate for those involved or interested in the specific modules or aspects of an ESG.

Section 11.4 comprises a list, without comments, of additional readings in ESG-related areas that provide either deeper analysis of material or alternative approaches to modeling. The final section, Section 11.5, offers brief suggestions for where future research may likely be found as it emerges.

Chapter 1: What Is an Economic Scenario Generator?

The chapter introduces the basic concepts, key considerations, and general characteristics associated with economic scenario generators (ESGs). Although many of the items treated here will be explored in more detail in later chapters, the discussion in this first chapter is relatively nontechnical and nonquantitative. Chapter 1 provides context and an overview of ESGs; it can also serve as a stand-alone introduction to ESGs—their characteristics, benefits, and limitations—for those unfamiliar with the topic.

1.1 The Concept of an Economic Scenario Generator

An ESG is nicely self-descriptive: it generates economic scenarios. More specifically, an ESG is a computer-based model that provides many simulated examples of possible future values of various economic and financial variables. These scenarios, along with analysis of the stochastic distribution of scenario outcomes, illuminate the nature of risk elements within the economy that drive financial variability. As such, an ESG can provide insights into the relative advantages and disadvantages of alternative operating and strategic decisions.

Since the term *economic scenario generator* is occasionally used differently by different people, it is important up front to specify what is meant by an ESG in this document. The main ambiguity comes with respect to the scope of the model that is being referred to. In this guide, we will use the narrow (but technically accurate) definition of an ESG being a model that simulates possible future economic scenarios. Because an ESG is typically used in combination with additional model components—components that use the economic scenarios as inputs, and then calculate those items of organizational interest (e.g., future firm value) for each simulated scenario under a variety of operational and strategic assumptions—it is not uncommon to hear the *entire* model, not just the economic-scenario-generating portion, referred to as an ESG. We will not do that here; in this document, *ESG* refers just to the portion of the overall model dealing with the projection of economic scenarios.

1.1.1. Value and Uses of an ESG

The value of an ESG is in its ability to simulate and project economic scenarios in a structured and rigorous way, as well as to explore distributions of outcomes that can illuminate risk. This is necessary because financial and economic variables are stochastic—they change over time, and in a largely unpredictable way. Only by considering a vast number of possible future economic scenarios—such as are generated by a good ESG—can we adequately gauge the breadth and relative likelihood of different possible future financial and economic conditions.

Despite the uncertainty inherent in future economic and financial conditions, organizations still must try to determine the best approach to operating in that future, a task that entails quantifying the risks associated with alternative approaches. In other words, they must make decisions in the presence of uncertainty. An ESG is a tool to help provide a framework for quantifying uncertainty—by structuring and quantifying that uncertainty in the form of probability distributions and stochastic processes. Describing this process is the primary subject matter of this guide.

An understanding of ESGs is not complete without an understanding of the broader models incorporated with the ESG, as well as their purposes. Here are some of the uses to which ESGs are put by property/casualty (P/C) insurers:

- Strategic and operational decision-making—e.g., underwriting decisions
- Solvency and solidity analysis
- Capital modeling
- Evaluating market risk
- Asset allocation
- Risk management—e.g., ceded reinsurance decisions

Basically, any decision involving the operation or strategic planning of a P/C insurer can benefit from—or even needs—the projection and analysis of effects of alternative decisions across an entire range of potential future economic and financial conditions.

1.1.2. Parts of an ESG

Like any model, an ESG model can be characterized by three major parts: input, output, and the calculations that go on in between.

- **Inputs** to an ESG: The users of an ESG will either take responsibility for the inputs to the system, focusing on those that are relevant to their situation and analysis, or rely on a third-party ESG vendor to undertake the task on their behalf. Inputs generally take the form of information regarding the economic and financial variables the ESG is simulating. Typically, that information involves selected parameters for each variable in the ESG. As time passes, newly emerged real-world data become available and are incorporated into an analysis that determines whether the parameter values already populating the model need to be changed.

- **Outputs** from an ESG: Descriptively, the basic output from an ESG is straightforward: each simulation includes a stochastically generated time series for each economic and financial variable contemplated by the model across a prespecified period of time (five years, 10 years, 30 years, 50 years, etc., depending upon the model and its uses). The model may display aggregations of these time series across all simulations—e.g., the average simulated three-month Treasury bill rate five years from now—but the real value of the ESG is in the individual simulations of each time series, as well as the distribution of those simulations across the range of future economic and financial conditions. This is because, as mentioned earlier, ESG output serves as the input to other sections of an overall model. Each simulated set of future conditions provides the basis for determining the company's condition and operations if that set of future conditions were to occur.

- **Calculations** within an ESG: Within an ESG, specific calculations take on several forms:
    - *Random number generation.* This allows the generation of random values from a prespecified probability distribution.
    - *Simulation of key variables.* One or more variables—often the risk-free interest rate, but in some models, inflation or another variable—form the foundation for the overall economic scenario.
    - *Correlation approach.* Based on cascade structure and other intervariable dependencies, other variables are simulated in a manner such that they are consistent (but still stochastic) with the key variable(s).

It is important to reemphasize that while ESGs are interesting in and of themselves, and we are discussing them in their narrow, literally defined sense, ESGs are generally one component of a larger risk management framework. In particular, the output from an ESG—simulations of possible future economic and financial conditions—is itself an input to a model that addresses the analytical issue being investigated.

1.1.3. Characteristics of an ESG

ESGs can take on a variety of specific structures but typically have several general characteristics in common:

- All critical economic and financial variables are simulated. The number and identity of the variables generated by an ESG depend both on the economic "depth" and detail desired and the purposes for which the ESG will be used.
- There is typically one economic variable (but possibly more than one) that serves as the "driver" of the other variables being generated in a scenario. Factors driving risk-free interest rates (for example, the three-month U.S. Treasury bill interest rate, or in multifactor models, the first n factors) probably constitute the most common driver variable.
- Projections of economic and financial variables are developed on a holistic basis. The values of different variables will be consistently generated in any given simulation, due to prespecified correlations and cascading dynamic variable interrelationships.
- All important asset classes are covered. For example, an entire panel of interest rates should be generated in order for bonds of different types to be appropriately modeled; equity returns must be generated to account for stock holdings; etc.
- Parameter values can be input and updated by the user. As economic conditions change and additional data emerge over time, users will likely wish to periodically reanalyze, recalibrate, and re-parameterize the model.

1.2 Component Modules of an Economic Scenario Generator

An ESG typically comprises several interacting modules, although the specific nature of the modules and their interrelationships may vary from one ESG to another. An illustrative example could be an ESG that has a "module" for each economic or financial variable for which the ESG generates scenarios. Although the specific items generated vary among ESGs, a representative list might include the following variables:

- Short-term, risk-free interest rates
  - E.g., returns on three-month U.S. Treasury bills
- Longer-term interest rates
  - Possibly by directly generating a yield curve
  - Possibly by adding a term premium to short-term rates
- Non-risk-free interest rates
  - E.g., municipal bonds, corporate bonds
  - Generated directly or by adding a default premium to risk-free rates
  - Default rates
- Equity market returns
  - Possibly in multiple categories, e.g., large stocks versus small stocks
  - Possibly generated according to different stock return "regimes"
- Dividend yields
- Mortgage-backed securities
  - E.g., collateralized mortgage obligations
  - Assumes a user-prescribed pattern for mortgage prepayments
- Possibly other assets
  - Derivative securities
  - Preferred stock
  - Other debt-backed securities
- Inflation
  - Possibly multiple indices or sub-indices such as wage, construction, or medical
- Other economic variables
  - Unemployment
  - Foreign exchange rates and relationships

- Gross domestic product or similar measure of output or economic health

One other component of an ESG deserves important mention: the *random number generator*. Basically, the random number generator allows for the generation of stochastically simulated random values from probability distributions. This component may be centralized and used by each of the variable components.

Two aspects of modeling the economic and financial variables in the modules mentioned above are critical for an ESG to generate reasonable scenarios of future values. The first is the *parameterization and calibration* of the variables; this involves identifying relevant and appropriate historical data and employing statistical and econometric techniques to determine specific parameters to enter into the modeling equations. While much of this work can be fairly technical, the effectiveness and appropriateness of the parameterization and calibration can often be at least partially assessed by whether the end result passes a relatively nontechnical "eye test." For example, if a short-term risk-free interest rate module is generating future interest rate patterns that bear no resemblance to any historic patterns or values, the parameterization (or the model itself) may be inappropriate. On the other hand, if patterns and values that differ significantly from those that have occurred historically are *possible*, one does not necessarily want a module that just reproduces history.

A great deal of art and judgment is involved in the building and parameterizing of economic and financial variable modules. Parameterization and calibration are discussed in more detail later in this *User's Guide,* particularly in Chapters 4, 7, and 9.

The second critical aspect of modeling the economic and financial variables involves the *interrelationships and correlations* among the variables. Building in the proper relationships between variables is essential for generating future scenarios that are internally consistent. A classic example is the relationship between inflation and interest rates. Would it make sense for an ESG to regularly generate scenarios that projected *both* very high inflation *and* very low interest rates? Probably not, and thus the interrelationship between the interest rate and inflation modules should reflect that likelihood (or unlikelihood) when building an ESG.

For the above example, and as mentioned earlier, very often an ESG will be designed so that one of the two variables—interest rate or inflation—is generated and then the other variable is generated assuming a dependent relationship on the first. But which variable comes "first"? This might be determined by an economic or philosophical decision or by a modeling decision, or by both. Again, designing and creating an ESG involves much art and judgment. Module interrelationships are discussed more in Section 1.3 below.

1.3 Relationship and Logic between the Component Modules

In the prior section, an example—interest rates and inflation—was presented to demonstrate the need for an ESG to contain appropriately structured relationships between modules when the variables generated by those modules are related (and that relationship could either be exhibited through empirical analysis of historical data or hypothesized intuitively or philosophically). We can model the interrelationships between economic and financial variables in three common ways. While the three approaches do each have distinctive elements, there is some overlap between them.

1.3.1. Cascade Structures

One definition of a *cascade* is a waterfall where the water falls in stages down rocky terrain. Another definition relates to the passing down of knowledge, say from one person or group or generation to another. Similarly, a cascade structure within an ESG involves the passing "downward" of information from one module to the next. For example, the level of the risk-free interest rate simulated for a future year may be passed to the simulated interest rate for the same future year. In turn, the simulated inflation rate, along with the simulated risk-free interest rate, may then be passed down to the next module (say, the equity return module). In a cascade structure, the simulated value of the "higher" variable becomes an input in simulating the "lower" variable.

This is a clean and methodical approach to designing an ESG. Nevertheless, such a model will require the designer to make several assumptions and decisions. For example:

- Are earlier values of the variable that is being simulated part of the information that provides the basis for simulating the variable? In the example above, the simulation of the inflation rate for a future year is a function of the risk-free short-term interest rate for that same future year; is it also a function of historical inflation rates up to that point?
- Are earlier values of the "higher" variables part of the information that provides the basis for simulating a variable? In the example above, in addition to the simulated future value of the risk-free, short-term interest rate, are historical values of that rate also a basis for determining the simulated future inflation rate?

1.3.2. Correlation Mechanisms

Sometimes, the simulated values of certain variables are determined "concurrently" or "contemporaneously"—i.e., they are conceptually and schematically "side by side," rather than vertically sequential in a downward cascade pattern. In that case, simple correlation metrics can be built into the modules that allow the generation of such variables to be simultaneous but with a prespecified degree of correlation between the variables. The specified correlation need not have a constant magnitude. Like almost any factor in an ESG model, an "expected" or "average" level of correlation between two simulated variables can be specified, along with a variability assumption (e.g., a standard deviation of the correlation value), allowing the degree of correlation in any simulated future year to be randomized within certain parameters.

1.3.3. Direct Linkages

Sometimes, a system of direct linkages reflecting the relationships can be imposed into the model structurally, as, for example, when there are multiple national economies involved. In the multination example, foreign exchange rates and relationships might serve as the direct linkages. A simple version of a direct linkage between two variables might be a simple linear model—perhaps with a stochastic error term thrown in—that could be parameterized by performing a linear regression on the variables' historic time series.

1.4 Risk-Neutral and Real-World Economic Scenarios

The concepts and mathematics of "risk-neutral" considerations can be a confusing aspect of advanced studies in financial economics. This is largely because, by its very nature, a risk-neutral framework does not have an intuitively comprehensible and transparent "real-world" manifestation. We are accustomed to dealing with people who are generally risk averse, and so to consider a world where everything and everyone is risk neutral, and to fully understand the implications of such an assumption, is difficult for us. Nevertheless, modeling within a risk-neutral framework is an important aspect of financial scenario generation—and regulators sometimes require it.

For ESG applications, then, modeling typically utilizes both risk-neutral and real-world frameworks. For example, the real-world price for a financial option is the expected value of the option's possible future values (or potential cash flows) under a general assumption that all investors are risk neutral—i.e., under a risk-neutral process. Similarly, bond prices are typically modeled as consistent with the fundamental theorem of asset pricing, which says that the price of an asset or a liability is the expected value, under the risk-neutral or equivalent martingale probability measure, of the discounted future cash flows.

An important application of this is a **market-consistent (risk-neutral) valuation framework,** which certain regulatory authorities require for the valuation of insurance liabilities. The market-consistent perspective indicates that the value of a liability should be quantified consistently with market prices in the sense that if an insurer sought to sell or transfer its liabilities, that would be the cost the market would charge to do so. Thus, a risk-neutral or market-consistent framework is one that takes into account (or is able to "reproduce") all market prices.

Overall, then, both the risk-neutral and real-world frameworks have a place in the ESG environment.

1.5 Approaches to the Stochastic Architecture of an ESG

This section discusses the architecture of an ESG and explores the implications of different approaches to the ESG's stochastic design. In general, when modeling economic and financial variables, one can use numerous approaches—and most of them have very strong advocates (and often equally strong detractors). What's important to remember is that each approach has value and usefulness in certain situations. The *goals* and *objectives* of the analysis are what determine which approach is most appropriate. Even a difference as simple as short-term versus long-term projections may dictate the use of one approach over another.

1.5.1. Discrete-Time versus Continuous-Time Models

Perhaps the most basic distinction between models is how time is treated: as a discrete or a continuous variable. Whereas a discrete-time framework is generally easier to conceptualize and understand, a continuous-time framework often leads to more convenient mathematics, sometimes including closed-form solutions to certain formulas and relationships.

- A **discrete-time** framework involves variables taking on values at specific and separate points in time. Those values stay the same for a certain interval of time, and only (potentially) change at the next specific point in time. Thus, a graph of the variable value over time is a step function that jumps from one value to the next at the discrete time interval. The length of the discrete time interval depends upon the availability of data and the nature of the application. A financial model that supports short-term trading of common stocks may look at stock values every minute (or, conceivably, even more frequently than that). A long-term economic or financial model may involve looking at variable values every month or every year. In all of these cases, the discrete-time framework is like taking a picture at regular time intervals: each discrete point in time represents a snapshot of the variables involved.

- A **continuous-time** framework is one where the variable values apply not to a finite interval of time but rather to an infinitesimally small instant of time. Thus, the values can potentially change continuously through time. In a sense, this framework is a subset of the discrete-time framework, where the discrete time interval shrinks to zero. A true continuous-time chart would be one that has no discrete jumps: it can be drawn without removing the pencil from the paper, because a small enough interval of time can always be found to provide a variable value of any desired closeness to the previous one.

For ESGs, continuous-time models are often used for the conceptual framework and mathematical underpinnings of the simulations of variables. However, technically speaking, in the real world, any measurement of economic and financial values will of necessity involve a discrete-time framework—even if measurements of variable values occur as fast as physically possible.

The relationship between discrete- and continuous-time models is close and convenient, although not without noteworthy differences. For example, using a continuous-time framework allows flexibility in selecting the length of the discrete-time interval, which can even differ for data analysis and simulation purposes. On the other hand, an example of a subtle distinction is that assuming a process is Markovian in one framework may not necessarily imply that it is also Markovian in the other framework.

1.5.2. Econometric Models

Econometric models are perhaps most often designed and employed with the objective to better understand the workings of the economy, but practitioners also often use them for predictive purposes. Econometric models use historical relationships between economic variables to estimate possible future economic conditions. Generally, these historical and future relationships are considered within a specifically defined economic modeling structure—hence, they are sometimes referred to as *structural models.* Such models are *causal* models in the sense that they assume linkages between the cause-and-effect relationships of different economic variables operating in an economy.

1.5.3. Macrofinance Models

A common dichotomy in the field of academic economics is that between microeconomics and macroeconomics. *Microeconomics* ("micro" for short) studies economics and behavior at the level of an individual person or organization. *Macroeconomics* ("macro") deals with economics on a global, national, or general economic level. As such, macro is involved with many of the broad economic metrics that we often hear about in the news: interest rates, employment, productivity, business cycles, etc.

As the name suggests, macrofinance examines the relationships between economic variables and asset prices. In particular, a macrofinance model is one that seeks to understand how changes in economic values and asset prices are related.

Extensions of modeling applications may reflect specific impacts of economic outcomes or behavioral responses of the organization to economic outcomes generated by the ESG.

1.5.4. Arbitrage-Free versus Equilibrium Models

The modeling of the term structure of interest rates is often one of the key elements of an ESG, and one of the principal distinctions between different philosophies of and approaches to modeling interest rates is between "arbitrage-free" and "equilibrium" models. Here we present an overview of those two approaches, and an indication as to the kinds of analytical tasks for which each of the approaches might be more appropriate. We discuss the issue in more detail in Chapter 6.

An *arbitrage-free model* is one where the relationships between economic and financial values do not allow for the possibility of arbitrage. This is done by recognizing the relationships between the prices of different instruments. For example, consider two portfolios, one that includes just an option on a share of stock and the other that includes both a share of the stock and cash (borrowing or lending at the risk-free rate). These two portfolios can be set up so that they have the same value in all possible states—i.e., the portfolios "replicate" one another. If two portfolios will always have the same price at a point in the future, then they must have the same price now. Thus, there is a mathematically expressible relationship between the two portfolios, and thus between the option and the underlying stock (plus cash). There is no possible arbitrage derivable between the portfolios.

An arbitrage-free interest rate model takes this idea and creates an interest rate process that does not allow for any arbitrage opportunities. Because of their nature and characteristics (more detail is provided in Chapter 6), arbitrage-free models are most commonly applied for short-term applications and for pricing derivatives. This type of model should also be considered the preferred choice for investment management applications where the presence of arbitrage opportunities within a model may cause overallocation to a particular asset class.

An *equilibrium model* of interest rates is one that begins with broad economic assumptions and then specifies the process as one that balances the economy (e.g., supply and demand). Such models tend to be better for longer-term horizons and applications of ESGs generating as they do asymptotically stable simulated economies.

1.6 Sector and Geographic Detail in ESG Applications

We can measure the breadth of an ESG in several ways: for example, the number of economic and financial variables it simulates or the range of different mathematical models that can be deployed in generating scenarios. Two other factors are also worth mentioning: the number of asset sectors that are represented, and the range and level of detail associated with economic and financial variables across different geographic and/or jurisdictional areas.

Depending upon the company and its ESG applications, there may be a higher or lower requirement for granularity in the ESG. A large part of the design process for an ESG involves decision-making regarding just how detailed the model should be. In some cases, those decisions are based not only on the specific needs of the company but also on the availability of data. Here are some examples of potential breakdowns of an ESG into different sectors or areas:

- **Asset classes:** modeling financial return data for different asset categories. For example, an ESG will certainly have modules for economic and financial variables associated with both the company's bond and stock portfolios, but should returns and valuations of different bond and equity sectors also be incorporated in the model?
- **Economic variables:** modeling certain economic variables, such as employment or inflation, for subsets of the economy. For example, are inflation rates needed just at the overall consumer price index level, or also for separate sectors such as wage, manufacturing, construction, etc.?
- **Geographical area:** modeling economic and financial variables by geographical area. For example, a company with significant business in different countries may very well want or need to analyze its operations in geographical subsets, which would likely necessitate having knowledge of the differing economic and financial conditions in each area.
- **Regulatory or jurisdictional area:** modeling variables according to regulatory spheres of influence. For example, regulatory rules might require that operations and financial reporting be segregated by jurisdictional location. There may be some overlap of this with the above "geographical area" considerations.

The more sectors and areas programmed into the ESG, the more comprehensive and responsive to detail the model will be. However, the ESG will also be more complicated to build and parameterize, and it may take significantly longer to run (see Section 1.8 below).

1.7 Benefits of ESGs versus Deterministic Economic Scenarios

Within the actuarial, insurance, and risk management communities, the movement from deterministic to stochastic models has been the result of long-term evolution. Over the decades, researchers and analysts have recognized that economic and financial quantities are not temporally constant and that markets have become increasingly volatile. Indeed, historically, there have been months when the interest rate on three-month U.S. Treasury bills changed by one or two hundred (or even more) basis points.

Moreover, it is not merely that economic and financial values can change over time—it's that they often do so in a random way. Change in and of itself is not necessarily an analytical problem if it is somehow regular or predictable, but when a time series appears to change randomly, the implications are difficult to measure and evaluate. This is why the ESG approach of simulating a range of possible economic scenarios is so valuable: we see the impact on an organization's financial condition and operations across a large number of possible futures, not just based on one deterministic, prespecified scenario.

However, situations still occur where actuaries may use deterministic paths for future values of economic variables; indeed, some regulatory jurisdictions require consideration of such scenarios. For example, New York Regulation 126—often referred to as the New York 7—promulgates seven deterministic interest rate scenarios that insurers must consider and against which evaluate their financial solidity via cash flow testing.

Such scenario testing (sometimes called stress testing) can be very useful for companies and regulators to better understand whether certain general kinds of future economic scenarios might lead to potential cash flow or valuation difficulties. However, there is one thing this type of scenario testing does not address: the relative *likelihood* of such a deterministic scenario actually occurring. Another drawback is that, for any finite-numbered group of test scenarios, the prespecified group of deterministic scenarios represents only a minute fraction of the possible patterns of economic values that *could* happen.

Compared with deterministic economic scenarios, stochastically generating economic scenarios has several crucial advantages:

- **Scenario probabilities:** As long as the probability distributions and parameters associated with the various economic variables are well specified and reasonable, random simulations generated from those distributions should be probabilistically spread across all simulated scenarios. In other words, inherently more likely values of variables should occur relatively more frequently across all simulated scenarios and inherently less likely values should occur relatively less frequently.
- **Scenario ranges:** Because each future variable value is being simulated many thousands (perhaps even millions) of times, there is a broad range of scenarios available to be tested and upon which to evaluate an organization's operations. Some scenarios will have paths of economic variables with values only slightly different from one another, while other

scenarios will reflect vast potential differences in future values. This spread of scenarios is much more robust than a finite number of deterministic scenarios could possibly be.

- **Scenario complexity:** Because of an ESG's various economic variable modules and their programmed interrelationships, a natural outcome of ESG scenario generation is scenarios with variables whose values have complex but reasonable relationships with one another. This capability allows an ESG to naturally produce scenarios that would be difficult to create artificially and deterministically. For example, the effect of multiple economic or financial variances can produce compound outcomes that may not be anticipated in a deterministic analysis of individual potential outcomes.

1.8 Limitations of an ESG in Modeling an Economy

ESGs have important benefits but, as with any type of model, they also have certain limitations, some of which apply to models generally, while others are specific to ESGs. The limitations do not invalidate the use of ESGs, nor should they dissuade one from using an ESG for appropriate purposes. However, those closely involved with such models should understand and appreciate these issues so that they can make appropriate provisions to overcome or mediate the issues and can work with and interpret the results of the model in the proper context.

- **Black box:** Like many sophisticated and complex models that are built, run, and maintained by a small subset of a larger audience, ESGs have a *black box* problem. Complete understanding of their construction and operation is limited and not easily visible without extensive investigation and deconstruction. A third party can see the input and the output of an ESG, but not what goes on in between. The key response to the black box issue involves *transparency*, *auditability*, and *documentation*. The model should be *tested*, *peer reviewed*, and *validated* when first built, and it should be periodically audited. The relationship between input and output should also be examined regularly for reasonableness.

- **Model risk:** Model risk involves the potential for mistakes or miscalculations in the construction, use, or interpretation of an ESG. Such errors could result from design flaws (e.g., a misspecification in how a variable value is generated, or in the linkage between variables), misuse (e.g., employing bad or inadequate data), or misinterpretation of results (e.g., incorrect or poorly structured output). The process for managing model risk is similar to the response mentioned above for the black box perspective: *testing*, *peer review*, and *validation*.

- **Maintenance:** All models require maintenance, although possibly for different reasons. For an ESG, the primary maintenance involves updating the model, especially with respect to its internal parameters. The emergence of real-world economic and financial variables and conditions needs to be monitored, and the ESG's internal parameters must be changed when conditions warrant. One approach to updating is to adjust parameter values one at a time, as the data dictate, leaving other parameter values as they were.

Another approach is to periodically recalculate and reinstall the entire panel of parameters in the model, based on the most recent economic and financial data. Under either approach, users should monitor the output and results for drastic changes due to updating; while changes in simulated variable patterns are expected and appropriate over time as parameter values are updated, drastic changes in output will need to be justified and may be undesirable. A secondary reason for regular ESG maintenance is to ensure that the model continues to meet the needs of its stakeholders, is generating the appropriate variables, and is producing reasonable results and output.

- **Complexity and resource requirements:** Software that attempts to model the future evolution of the economy and financial markets is an inherently challenging and complex undertaking. At some point, the complexity and size of the model may lead us to question whether an additional, marginal improvement or upgrade will provide a sufficient improvement in the modeling capability to make it worthwhile. ESGs have the potential to get too big, requiring enormous and possibly unreasonable resources to build and maintain, and also producing overlong running times. Although computers continue to become faster and more efficient, it is still necessary to plan and design an ESG with cost–benefit considerations in mind. Such considerations also apply to the amount of data and technical expertise required to run the model.

- **Noise, instability, and convergence:** An ESG attempts to model the future based largely on historical economic and financial conditions and patterns. Several issues need to be recognized:

    - A historical data set makes up a sample of what *might have* happened. The historical data may be representative of the average underlying tendencies in the economy at that time, or some of the data points may have been outliers among all the potential economic outcomes. It is important to keep in mind that historical data, which serve as the basis for parameterization of an ESG, are noisy. Historical data may also guide a modeler to an incorrect view about the future, due for instance to a structural change in a particular market.

    - This inherent noise in the data means that the ESG parameters, and thus the results, output, and the model itself, should be checked for stability.

    - Noise is found not only in historical data but also in the simulated output. The more complicated the model, the more potential there is for noise, and the longer (i.e., more simulations) it will take for the simulated values to converge. There are a number of rules of thumb one can use to suggest how many iterations a simulation requires before convergence of results, and there are also approaches to the simulation itself that can help with this issue—for example, employing quasi-random sequences rather than simple Monte Carlo sampling.

1.9 Guidance on the Use of ESGs

Several principles should guide both the design and use of an ESG:

- **Objective:** The objective and purpose of the analysis to be undertaken with an ESG should dictate the techniques and modeling formulas used. For example, if the purpose for building an ESG is to have a basis for making short-term financial trading decisions, it may make sense to consider employing a no-arbitrage model for generating interest rates. On the other hand, if the ESG is supporting long-term, multiyear operational and strategic decision-making, an equilibrium model of interest rates may be better.

- **Transparency and documentation:** Documentation is "perhaps the most important pragmatic issue for modelers" (Sterman 2000, pp. 865). Documentation, in addition to providing guideposts to the model programmers as they write the code, is the basis for the transparency necessary for the users and observers of a model to understand its operation. Thus, transparency and documentation in an ESG ideally promote the understanding of at least the following underlying aspects of the model:
    - Intent and goals
    - Assumptions
    - Variables and their interrelationships
    - Data employed
    - Output and its uses

- **Parsimony:** Conceptually, a model can always perform better—i.e., reduce the model residuals relative to the data—by adding another parameter or explanatory variable. However, having too many explanatory variables can cause problems—for example, while the fit to historical data might be good, an overparameterized model is unlikely to be much good for future predictive purposes. In addition, a model with an extremely large number of explanatory variables would be inconvenient (at best) to actually use.

  Instead, the goal in modeling is generally not to get the best fit to the current data at all costs, but rather to minimize the number of explanatory variables while maximizing the goodness of fit. In other words, the ideal is to use no more explanatory variables than necessary to achieve a satisfactory fit. This is the principle of "parsimony."

  Therefore, parsimony involves a trade-off: is the improvement in model fit large enough to justify the addition of another explanatory variable? Making such a decision involves a degree of judgment. It also requires ways of measuring the improvement (discussed further in Chapter 5); some of the techniques used to compare the parsimony of models include the Akaike information criterion and the Bayesian information criterion.

- **Time and resources:** To say that computing power and capabilities have improved over the last two or three decades would be a monumental understatement. One aspect of this multifaceted improvement is time: processing speeds have increased to the point where a program or simulation that might have required an overnight, or even a several-day, run a decade ago can often now be accomplished in minutes. This greater speed provides an important advantage in the use of sophisticated models such as ESGs: the ability to interact with the model in an almost real-time framework, changing assumptions or inputs as they emerge (or as they are dreamed up by users), and getting virtually immediate feedback on the implications of those changes. So, certainly, the time needed to run an ESG model has drastically decreased from a computer-usage perspective.

Interestingly, though, the necessary time and resources associated with running and managing an ESG *in its totality* have not uniformly decreased. As technical expertise and capabilities have improved, so has what is asked from a model. We are asking more and more from ESGs, both in breadth and depth. Compare an ESG from 10 or more years ago with one that is being designed today, and you are likely to see one or more of the following enhancements:

- More economic and financial variables being projected
- More sophisticated economic and financial theories or paradigms being used as the basis for modeling certain variables
- Greater depth and breadth of output, to address a wider variety of operational and strategic issues that ESGs are being used to inform
- More complex and sophisticated interrelationships between variables

Perhaps the most significant enhancement of all is the greater availability of more, better, and more detailed data, along with greater resources (both human and software) with which to analyze those data.

So although computing speed has significantly shortened the amount of time needed to run an ESG model, the greater scope, purview, and sophistication of many ESGs have resulted in more resources devoted to the data (collecting, managing, analyzing, calibrating) and to the general maintenance and operation of the ESG and the systems to which those data are attached.

1.10 Summary

An economic scenario generator is nicely self-descriptive: it generates economic scenarios. More specifically, an ESG is a computer-based model that provides many simulated examples of possible future values of various economic and financial variables. Those scenarios, along with analysis of the stochastic distribution of scenario outcomes, illuminate the nature of risk elements within the economy that drive financial variability. As such, an ESG can provide insights into the relative advantages and disadvantages of alternative operating and strategic decisions.

An ESG is typically used in combination with other models—components that use the economic scenarios as inputs and then calculate those items of organizational interest. The value of an ESG is in the ability to simulate and project economic scenarios in a structured and rigorous way. This is necessary because financial and economic variables are stochastic—they change over time in a largely unpredictable way.

An ESG typically comprises several interacting modules, although the specific nature of the modules and their interrelationships may vary from one ESG to another. Like any model, an ESG model can be characterized by three major parts: input, output, and the calculations that go on in between. An ESG typically simulates all relevant economic and financial variables, but one variable tends to serve as a driver of the other variables being generated in a scenario. Projections of variables are developed on a holistic basis, and asset classes are covered. Finally, the user can update parameter values, and the ESG can be assessed and validated.

Two critical aspects of modeling the financial and economic variables in the modules of the ESG include the parameterization and calibration of the variables and the inclusion of proper relationships of correlation and other interrelationships between the variables. The interrelationships may be developed by a cascade structure, correlation mechanisms, and direct linkages.

For ESG applications, modeling can follow either risk-neutral or real-world approaches. Some regulatory authorities require risk-neutral (or market-consistent) frameworks for valuation of insurance liabilities. Real-world modeling is appropriate when projecting future values of economic and financial variables. Modeling can also follow discrete-time and continuous-time mathematics. Generally, continuous-time modeling leads to more convenient mathematics.

Analytical tasks may also distinguish between arbitrage-free and equilibrium models. Arbitrage-free requires that the relationships between economic and financial values do not allow for the possibility of arbitrage. An equilibrium model specifies that the interest process balances supply and demand and is often better for longer-term horizons.

Compared with deterministic economic scenarios, econometric models, and macrofinance models, an ESG simulation can provide a better view of scenario probabilities, a broader range of scenario outcomes, and greater complexity of scenarios. However, an ESG can appear as a black box, may contain significant model risk, and can require significant resources for maintenance and development. With increased computer resources and better and more detailed data, the greater scope and sophistication of ESGs has resulted in more resources devoted to data and to the general maintenance and operation of the ESG. Thus, it is critical that the design of an ESG carefully consider the objective of the analysis to be undertaken, with sufficient attention being paid to transparency and documentation.

References

Sterman, J. D., 2000. *Business Dynamics: Systems Thinking and Modeling for a Complex World*, Boston: McGraw-Hill, 2000.

Chapter 2: Applications of Economic Scenario Generators

In the previous chapter we discussed the basic concepts, key considerations, and general characteristics associated with ESGs including the interrelated components, benefits, and limitations. In this chapter we review the various insurance applications that require or can make use of an ESG. The purpose of this chapter is to address the following questions:

- How are ESGs used to identify and manage risks to an insurance company?
- How are ESGs used within an insurance company's business and reporting processes?
- How are ESGs distinct from trading models and forecasting models?

The key aspects of any ESG application revolve around establishing a probability distribution of future economic-influenced results. An ESG application will generally place greater emphasis on the variability of results rather than simply the overall mean of those results—in order to provide greater insight and information in the context of economic risk management related to a company's operations.

This chapter provides several illustrative examples of how an ESG could be used in the context of insurance. For example, an ESG application could aid in developing strategies for bond holdings by risk and maturity or in evaluating alternative reinsurance strategies, as well as policy limits, on the insurance side.

For the purposes of distinguishing key features and uses of different ESG applications, we divide this chapter into different aspects of risk management (as set out in the next section) so as to provide context for how those aspects would be included and addressed in the various applications discussed in subsequent sections.

The final sections of the chapter include a comparison and contrast with some other modeling techniques as well as a discussion of some of the practical aspects to be considered in the context of ESG applications.

2.1 Risk Monitoring, Management, and Control

Broadly speaking ESGs are used to support one or more of the following key aspects of effective risk management:

- **Monitoring:** ongoing measurement and reporting of the various risks to which an insurance company is currently exposed (or considering) based on its business operations
- **Management:** establishing, updating, and executing risk response strategies on an ongoing basis
- **Control:** ensuring that the above aspects are comprehensive, depictive, and performed on a timely basis

It should be understood that the above categories are neither necessarily rigid nor nonoverlapping but are used here to delineate different aspects of the various purposes and results supported by the illustrative ESG applications discussed in this chapter. In practice, a given ESG application might support more than one aspect, and a company's risk management practices could use terminology different from the categories named above.

Most ESG applications, as discussed in subsequent sections, provide support with respect to economic risk in the production of results and analysis for *monitoring* and *management* purposes.

Control aspects relate to the governance of risk management and, in the context of ESGs, would relate to items such as breadth of economic factors (ensuring all relevant economic risks are included) as well as agreed-upon and documented practices for the setup and maintenance of ESGs.

Loosely speaking, monitoring relates to the *evaluation* of some or all of a company's financial positions as of a point in time, whereas management relates to assessing and establishing *strategies* a company would employ for dealing with risk. In practice, a given ESG application might reflect elements of both evaluation and strategy—such as an application to assess an existing or proposed investment strategy where some of the key metrics used to assess effectiveness involve future evaluations of the liabilities over time as they change in response to simulated economic conditions.

To better draw out some of the key elements and specific purposes of different ESG applications, the following discussion has been partitioned into *monitoring* and *managing* separately—knowing that a given application, while likely focused on one of those two facets, might include elements of the other in its overall results.

2.2 Applications of ESGs—Monitoring Risk

2.2.1 Measuring and Monitoring—Liability Risks

Liabilities in this context make up the obligations of an insurance company owing to the insurance and reinsurance contracts that it has entered into or that might be established in the future. Owing to these obligations, insurers will pay or receive monies subject to events that have already occurred or may occur in the future.

In estimating the value of current contracts and in establishing prices of future/renewal contracts, insurers need to make assumptions and estimations related to the timing and amounts of future payments. The evaluation of current contracts might be established using different methods and assumptions depending on which set of financial statements is relevant.

The applications used for such purposes are typically referred to as *reserving systems* (for current contracts) and *pricing systems* (for future/renewed contracts).

The amounts of such future claim payments, and, if applicable, their discounted present value, can vary subject to the risks modeled via an ESG. The following are some of the more common risk factors addressed using ESGs within a liability system:

Inflation

For property/casualty insurance, inflation often plays a significant role with respect to future payments. Some liability systems may incorporate an ESG to simply establish an average level of future inflation for reserving or pricing purposes, whereas other liability systems might generate a stochastic set of future payments using the ESG-generated inflation rates.

Unemployment Rates

Historically, insurance disability claims have displayed some correlation with respect to unemployment rates. Liability applications that incorporate this effect can make use of an ESG to establish a robust and reliable distribution of unemployment rates that can then be applied in estimating such claims on a forward-looking basis. In this context, the parameterization of the ESG will be critical to producing dependable and actionable results.

Risk-Free Interest Rates and Corporate Bond Spreads

While these economic risk factors do not typically directly influence the timing and amount of payments, they can affect the discounted present value of insurance payments—and therefore the aggregate liability of current contracts and even the prices established for new or renewed contracts.

In addition, interest rates and inflation rates would typically be materially correlated so that changes in interest rates would be correlated with changes in future claim payouts whether they are used for discounting or not.

Often these economic risk factors can be established as of a given date within a reserving or pricing system, based on then-current economic conditions, without any requirement for an ESG. However, for longer-term insurance products and/or applications addressing not just liabilities (as described in subsequent sections), the proper measuring and monitoring of liabilities may involve consideration of future economic conditions and their impact on either future liability reserves or investment activities.

2.2.2 Measuring and Monitoring—Investment Risks

We can use an ESG for two distinct aspects of the measurement and monitoring of an insurer's investment risks—most often, for assessing the effectiveness of an existing or proposed ongoing investment strategy and/or, sometimes, for the valuation of private or nonliquid investments. The latter aspect is not common but is described (in the next section) for the sake of completeness.

Investment Valuation (Measurement)

Most investments an insurer holds can be readily valued as at a certain date, making use, where appropriate, of publicly available information from the financial markets (even though the current *market value* of an investment may be just one component used in determining the investment's *book value*). For some investments, such as real estate or private equity, there may be no or limited publicly available information with which to reliably determine its market value—such investments might be repriced infrequently or only when they are bought or sold.

Depending on the nature of the investment, an ESG might play a role in the proper assessment of that investment in the absence of a publicly available value.

As an illustrative example, the value of a real estate investment could be determined (or simply tested for reasonableness) by taking into account a stochastic simulation of future rental prices and occupancy rates—which in turn are driven by economic risk factors such as inflation and gross domestic product (GDP).

Investment Strategy Assessment (Monitoring)

The ongoing measurement and monitoring of an insurer's established investment strategy is a critical function and one in which an ESG plays a major role. The ESG application used in this context—i.e., measurement and monitoring—is often the same system used for management and mitigation as detailed in Section 2.3.2, "Managing and Mitigating—Investment Risks."

An insurer should regularly monitor the economic risks associated with its current investments and how it plans to manage its investments going forward.

Within such an application, one may use the range of future simulated economic conditions to establish the corresponding range of future investment results. Such an assessment will demonstrate the likelihood and severity of overall economic risk and could also be used to delineate the relative impact of individual market risks such as interest movements and equity markets.

2.2.3 Measuring and Monitoring—Capital Risks

While insurance capital can be defined and established in various ways, for the purpose of this discussion we will define it simply as the excess of investments over liabilities. One of the primary purposes of capital is to serve as a buffer against extreme adverse events, allowing the insurer to continue satisfying its liability obligations following such an event. These are events that either exceed those anticipated and reflected in establishing loss reserves or were not anticipated at all—this discussion will concentrate on the former.

Capital risk measurement and monitoring often revolves around the consideration and establishment of risk tolerances (the maximum risk an organization is willing to take regarding each relevant risk) and risk limits (thresholds with which to monitor that the actual risk exposure stays within an organization's risk tolerance).

Insurance regulators and rating agencies have established methods for calculating a sufficient level of capital, taking into account the various risks to which an insurance company is exposed.

While many of these *risk-based capital* approaches are based on detailed formula-driven calculations that one can perform without the use of an ESG, increasingly rating agencies such as A.M. Best (Wong-Fupuy and McGuigan 2020) encourage the use of an ESG-driven *internal economic capital model*—where an insurer has established its own view on how to measure and assess capital adequacy—as distinct from one promulgated by a regulator or rating agency.

A common approach to establishing a sufficient level of capital relies on a *1-in-200-year* basis (Morin 2011) in which, without defining it rigorously, a company should have sufficient capital to absorb cumulative adverse events 99.5% of the time as measured over a one-year horizon. For any significant risk an insurer undertakes, such a calculation requires a stochastic assessment—including a robust and reliable *tail distribution* so as to reliably observe where the 99.5th percentile occurs.

For most property/casualty insurers, the tail distribution can be dominated by noneconomic exposures (windstorms and other catastrophes), but, depending on the actual insurance contracts and any associated reinsurance arrangements, economic risks will likely contribute a meaningful amount to the overall calculation of economic capital (investment events like the 2008 financial crisis can erode capital in a hurry).

An ESG is critical to properly assessing economic capital that includes economic risk. It is critical that the ESG in this context performs well over the calculation horizon (one year, as cited above, is most common) and that the tail of the distribution is reliable and robust. Economic event tails are often made up of simultaneous events (e.g., an equity market drop combined with widening corporate spreads or sudden jumps in [expected] inflation)—making the integrity of the overall integrated model a top criterion.

Many economic capital applications address measurement—where monitoring might be simply checking that the calculated economic capital does not exceed actual current capital. Monitoring could also include reconciling the movement in the level of capital from the prior period—allocating discrete movements by individual risk factors in a *waterfall chart* and using an ESG to isolate the contribution of each market-risk component.

A more ambitious economic capital application would address management and mitigation (of capital erosion) as discussed below in Section 2.3.3, "Managing and Mitigating—Capital Risks."

2.3 Applications of ESGs—Mitigating Risk

Traditionally, there are four methods for mitigating risks:

- **Avoidance:** excluding a particular risk—e.g., avoiding currency risks by using only the U.S. dollar in all contracts
- **Transfer:** assigning or selling some or all of the risk to another party—e.g., reinsurance
- **Mitigation:** approaches used to reduce the severity of possible outcomes—e.g., duration matching to dampen the effects of interest rate movements
- **Acceptance:** taking on risk where transfer or mitigation costs are prohibitive

The next sections focus primarily on the mitigation of economic risk—dealing with ESG applications used for that purpose. This section builds upon the measurement concepts discussed earlier, since, in order to assess mitigation strategies, measurement capability is a requirement.

2.3.1 Managing and Mitigating—Liability Risks

For property/casualty insurers, the dominant risk mitigation strategies are noneconomic, with reinsurance being a common technique.

Where economic risk mitigation strategies are employed (e.g., using inflation-linked bonds to mitigate inflation risk), they are generally embedded in the liability measurement and monitoring applications discussed previously.

2.3.2 Managing and Mitigating—Investment Risks

Mitigation of investment risk is the most common ESG-driven application for insurers. Sometimes referred to as *investment management systems* or *asset–liability management systems*, these systems (we will use the terminology *ALM application* here) are primarily designed to assess and establish investment guidelines with a view toward mitigating exposure to economic risks while increasing the likelihood of positive investment returns.

The distinction between an investments-only system and an ALM system is the degree to which liability considerations are reflected. Either approach can be effective, depending on the strategies to be assessed and to what extent liabilities have significant economic risk. Liability data can include future estimated payments and/or future estimated loss reserves, both of which can inform the effectiveness of any investment strategy.

An ALM application typically serves to assess and establish an insurer's investment *strategy*, comprising high-level guidelines and targets for investment managers in the *tactical* execution of that strategy on a day-to-day basis. ESGs are a fundamental requirement of an ALM system but are not necessarily needed for the execution of the buying and selling of specific securities (for more discussion on this, see Section 2.5, "Applications of ESGs—Comparing Related Applications").

Broadly speaking, an investment strategy can be assessed on an *economic basis* (also known as *market value basis)*, on an *accounting basis* (also known as *book value basis* or *regulatory basis*), or possibly both. It should be noted than accounting-basis analysis typically requires that the ALM system be capable of simulating results for individual investment holdings (often called security-level modeling) in order to properly capture and reflect their respective book values over time; this requirement (which increases both application complexity and run times) extends to the ESG, which must be capable of generating economic data for this purpose.

Irrespective of the basis (economic or accounting), the risk factors to be included in an ALM system—and the aspects to be reflected in their management/mitigation—would consist of some or all of the following:

<u>Risk-Free Interest Rates</u>

The value of fixed-income investments will rise or fall respectively in response to changes in risk-free interest rates. These events can have an impact on an insurer's investment returns and capital position. Discounted liabilities would likely move in tandem with investments—providing some natural mitigation, but not necessarily an exact offset to interest rate risk or any associated movements in inflation rates.

For a proper assessment of interest rate risk, the ESG within an ALM application must model the future distribution of risk-free interest rates in a reliable and robust fashion, ensuring sensible behavior along a particular scenario and overall across all scenarios.

To model and assess interest rate risk mitigation strategies, an ALM system should be able to provide data and metrics related to the future cash flows of investments and, where applicable, liabilities. This might be simple duration (or key-rate durations) by fixed-income investment type or, at the extreme, individual cash flow projections for each individual fixed-income security.

A classic mitigation strategy for interest rate risk involves using investments with a slightly shorter duration than the liabilities—so that investment returns will be close to liability accruals while reducing the downside due to increases in interest rates.

Yield Spreads

In order to realize enhanced investment return potential, an insurer's investment strategy will typically address the use of rated fixed-income securities where the value of such investments reflects a spread above risk-free interest rates, providing for higher but less certain future returns and cash flows. The value of such investments will move not only with risk-free interest movements but also with movements in the "spread" between risk-free rates and credit yields and will exhibit comparable impacts to future investment returns and capital levels (although historically spread risk is of a lower order than interest rate risk).

The ESG within an ALM application should provide proper spread dynamics across the range of eligible fixed-income investment types within an insurer's investment strategy, and the simulated spreads must be sensible (e.g., AAA spreads lower than AA spreads) and be realistic within and across investment types.

The mitigation of spread risk within an ALM application must be addressed within the context of the related cash flow uncertainty as described below. Essentially, an insurer takes on spread risk and cash flow uncertainty in order to enhance return potential—a trade-off that requires an ALM application to properly understand and consider.

Defaults and Prepayments

Defaults of corporate bonds typically result in an instantaneous loss in investment value and simultaneous early (or even delayed) payment of that lower value. Defaults will reduce both investment returns and capital.

Prepayment risk can be a predominant feature of any asset-backed (particularly mortgage-backed) investment. Prepayments consist of earlier payments of an investment's principal that do not necessarily erode total value but often occur during adverse economic environments (e.g., prepayments often occur when mortgage holders refinance at lower rates, leaving an insurer with extra cash that must then be invested in a low-yield environment). Convertible bonds (which the issuing company can convert into stock) also include aspects of prepayment risk.

The ESG within an ALM system must provide the data necessary to properly reflect and assess default and prepayment risk—the frequency and severity of such events should be consistent and plausible within all the economic conditions simulated. Prepayment risk can be particularly challenging owing to the variety of mechanisms found within asset-backed securities.

Both defaults and prepayments interact primarily with interest risk mitigation, since the uncertainty of cash flows must be adequately addressed for the purposes of duration- and cash-flow-matching strategies.

Transitions

The rating transition of a fixed-income investment seldom has an instantaneous impact on investment returns or capital levels, but transitions (often more frequent in high-spread environments) can increase the likelihood of subsequent cash flow uncertainty. Rating transition can affect the cash flow uncertainty aspects of an insurer's investment portfolio and affect the regulatory/required capital levels, which are typically ratings dependent.

The ESG within an ALM application should provide sufficient rating granularity and robust transition dynamics to properly assess an insurer's investment strategy.

Traditionally, transition risk is treated as a sort of early warning sign, and an insurer's investment strategy would provide guidance with respect to rebalancing the rating proportions of its investments—how often and under what circumstances. An ALM application should provide the capabilities allowing for alternate rebalancing strategies to be depicted and understood. Modeling such rebalancing strategies can be challenging, owing to other desired investment portfolio attributes (duration, prepayment exposure, etc.).

Liquidity Risk

Liquidity risk is the risk that a company becomes unable to meet short-term financial demands. This usually occurs due to a limited ability to convert a security or hard asset to cash without a loss of capital in the process. An ALM system that addresses liquidity risk might allow for aspects such as transaction costs, bid-ask spreads, and illiquid/private assets.

An insurer can assess liquidity risk via stress testing—incorporating scenarios in which cash outflows are significantly elevated and/or incorporating liquidity "triggers" in which economic events (e.g., widening credit spreads) lead to a reduction in realizable asset values.

Equity

Traditionally an insurer would use equity investments to enhance overall investment returns while taking on additional uncertainty with respect to those returns. The allocation to equity is typically significantly less than that of fixed-income investments. Historically, equity returns can be larger (both positive and negative) and more variable than fixed-income returns, but they can also be superior to comparable fixed-income investments on average and over the longer term. Regardless, equity risk can erode both investment returns and capital (usually with no natural liability offset) so care must be taken in striking a proper approach to equity within an insurer's investment strategy. The ESG within an ALM application must be able to model equity dynamics within the context of an insurer's investment guidelines, providing sufficient granularity within and across equity markets and consistent with the overall simulated economic conditions including correlation with interest rates and credit spreads.

Equity risk is more often limited rather than mitigated by using an ALM application to assess and establish strict limits with respect to equity risk exposure—perhaps even in response to simulated capital levels over time. Equity risk mitigation, to the extent it is implemented, is often addressed via equity derivatives where some downside risk is eliminated—either at a price or in giving up some upside risk (see "Other Aspects" below for more discussion on derivatives).

<u>Foreign Exchange Rates</u>

Foreign exchange rate risk can arise when either investments or liabilities are denominated in currencies other than an insurer's reporting currency. Changes in exchange rates can affect investment returns and capital levels—both positively and negatively—increasing the variability of an insurer's investment strategy.

The ESG within an ALM application that reflects more than one currency must model not only exchange rates among currencies but also the financial markets within each currency's economy (at least those portions of the financial market in which an insurer is invested)—and must do so on a simultaneous and consistent basis within and across economies. Note that a multi-economy ESG simulation must allow for the added complexity in which all economies "move together" in a sensible/plausible fashion and yet provide for instances in which some markets or even entire economies "move against the tide"—this presents both a modeling and parameterization challenge.

Similar to equity risk, foreign exchange risk is often limited rather than mitigated, although mitigation becomes increasingly necessary when the currency exposures are significant. Mitigation could involve the use of foreign exchange derivatives (to transfer the risk to a third party) or even investment segmentation in which liabilities of each currency are backed by distinct investment portfolios in that currency. Note that the latter approach would require an ALM application that supports multiportfolio investment strategies.

<u>Inflation, Unemployment, and GDP</u>

Some ALM applications (and by extension, their embedded ESG) include such risk factors as inflation, unemployment, and GDP, particularly when the liabilities are sensitive to one or more of those factors (as discussed above in Section 2.2.1, "Measuring and Monitoring—Liability Risks").

<u>Other Aspects</u>

Some ALM applications can include specific aspects necessary either to assess an insurer's investment strategy or to enable a particular mitigation approach to one or more of the risk factors discussed above. The following provides a sampling of such aspects and how they might have an impact on the required capabilities of an ESG:

- **Tax accounting:** Many ALM applications operate on a pretax basis, leaving the complexities of taxation to another application run by a separate team. For those ALM applications that establish after-tax analysis, the embedded ESG typically needs to support security-level modeling in order to support proper aggregate tax calculations.

- **Derivatives:** Some mitigation strategies employ the use of derivatives—commonly equity options, foreign exchange swaps, or even credit default swaps. For these purposes, an ESG will need to support the pricing of such derivatives—and not just their current prices, but the simulated future prices consistent with emerging simulated economic conditions.

2.3.3 Managing and Mitigating—Capital Risks

As described earlier in Section 2.2.3, "Measuring and Monitoring—Capital Risks," an insurer's capital can be measured on either a regulatory basis or an economic basis.

Those applications intended to aid in the management and mitigation of capital erosion risk are often extensions of the ALM applications discussed in the previous section, where an insurer's investment strategy includes aspects related to capital levels and capital protection. For example, an insurer may increase or decrease its equity exposure when regulatory capital levels rise and fall over time, and a full economic capital ALM application that can measure capital levels in simulated future economic conditions can support such an approach. In this case, an ESG will need to provide any economic data required for such calculations.

2.3.4 Nested Stochastics

A very sophisticated ALM application might calculate future levels of stochastically based capital, such as the 99.5th percentile approach described earlier. In this case, an ESG would need to support *nested stochastics* in which a full stochastic simulation is created for each future time period in which capital is to be calculated.

Such applications are not commonplace for property and casualty insurers and are described here only for completeness. Nested stochastics models require a large amount of computer processing power but are seeing some traction in life insurance companies, owing to the "embedded options" found in some life insurance products.

2.4 Applications of ESGs—Stress Testing

Stress testing is a technique in which one alters economic conditions or economic simulations in specifically defined ways to either measure an insurer's exposure to some set of risk factors or test the sensitivity of an application's results to its underlying assumptions.

Cash flow testing is good example of changing economic conditions in a defined way in order to demonstrate an insurer's capability to absorb interest rate movements in the future (Blanchard and Marchena 1995). Ideally, the ESG used for stochastic analysis should be able to generate, or take as input, predefined scenarios, allowing those results to be compared on a like-for-like basis with the results from a stochastic run of the same application. This type of deterministic stress testing is a good supplementary reporting technique in that the underlying assumptions can be readily understood and related back to the calculated results, even if the *probability* of a particular deterministic scenario may not be well defined or understood. Such tests can also aid in communicating the appropriateness and effectiveness of risk tolerances or limits.

An ESG could be calibration stress tested in order to understand the implications of using one calibration over an alternate calibration. For example, within an ALM application the level and speed of the mean reversion in risk-free interest rates could be altered so as to determine how robust a particular investment is with respect to the underlying interest rate dynamics: does the alternate calibration result in materially different results that change the conclusions of the effectiveness of the investment strategy? Or the relative merits of one strategy over another? This type of stress testing is a critical step in the construction and ongoing maintenance of an ESG within an ALM application—understanding the extent to which the results and the decisions derived are reliable.

2.5 Applications of ESGs—Comparing Related Applications

To extend the discussion of ESGs within ALM applications, recall that ALM applications are primarily designed to assess investment *strategies*—leaving the execution of those strategies to other teams and other applications. One of those other applications could be a *trading model* in which specific potential trades are analyzed in real time, based on current market conditions, to see not only how a particular trade would align with the established trading strategy but also how it might compare in this context with other possible trades. A trading model typically does not use an ESG in the execution of its functionality and instead takes on real-time data, near-term trends, and even long-term trends in preparing its metrics for analysis.

A trading model may also incorporate a *forecasting model* that generally provides a smaller number of scenarios relative to an ESG (sometimes simply expected, best case, and worst case). Forecast models typically take on an ongoing, frequently updated, current best estimate of emerging economic conditions—as contrasted with an ESG application where the calibration criteria are adjusted less frequently (often only to update them to current market conditions). Forecasting model applications can exist independent of trading models, but they seldom include an ESG as described in this paper. Forecasting models might differ from ESGs in their consideration/reflection of the future range of results: a forecasting model might focus on short-term volatility with respect to best- or worst-case scenarios, whereas an ESG might have a longer-term projection horizon in which its parameters have been estimated to reflect long-term volatility of economic conditions.

2.6 Applications of ESGs—Some Practical Aspects

During the implementation and maintenance of an ESG-driven application, a host of practical considerations will arise that we must take into account, particularly with respect to the ESG itself:

- **Breadth and scope:** Is the ESG able to provide all relevant data required by the application? Can the ESG be extended in light of new application features such as additional investment types? Can the ESG reflect the required granularity of aspects such as asset classes and security types to be modeled?

- **Performance:** Usually, the run time of most ESGs is a small proportion of the overall run time of its application, but consideration should be given to calibration turnaround times and calculation-intensive security-level modeling requirements (where applicable).
- **Granularity:** Do the ESG outputs align properly with the desired granularity of the application? For example, does the ESG provide only total investment returns, whereas the application calls for price and income returns separately?
- **Setup and maintenance:** Can the ESG be set up and maintained on an ongoing basis within the application with a reasonable amount of effort? What expertise is required to run and validate results? What tools does the ESG provide for such purposes?
- **ESG limitations:** Does the ESG have any limitations that would materially affect the depictive and robust performance of the application? We address specific limitation considerations within various other chapters in this paper—particularly Chapters 8 and 9.

2.7 Summary

The most common ESG-driven applications for property/casualty are ALM systems (used in assessing, establishing, and monitoring investment strategies) and economic capital systems (used to calculate and monitor economic capital).

ALM systems deal primarily with economic risk mitigation, in which the range of adverse economic events is narrowed or reduced while still maintaining a healthy likelihood of positive investment growth.

Economic capital systems typically focus on shorter time horizons and involve significantly more scenarios in order to establish reliable tail metrics.

Any ESG application will have some practical limitations based on its underlying ESG and any functionality (trading strategies, etc.) that the builder has provided or implemented in support of the application use case(s). Users of an ESG application must appreciate any such limits in order to appreciate how best to interpret and communicate results generated.

References

Blanchard, R. S., and E. P. Marchena, "A Decade of Cash Flow Testing—Some Lessons Learned," Casualty Actuarial Society Discussion Paper Program, 1995, https://www.casact.org/pubs/dpp/dpp95/95dpp023.pdf.

Morin, F., "ERM-2: Introduction to Economic Capital Modeling," a presentation at the 2011 Casualty Loss Reserve Seminar, Las Vegas, September 15, 2011, Towers Watson, https://www.casact.org/education/clrs/2011/handouts/ERM2-Morin.pdf.

Wong-Fupuy, C., and M. McGuigan, "Best's Credit Rating Methodology," A.M. Best Company, March 5, 2020, http://www3.ambest.com/ambv/ratingmethodology/OpenPDF.aspx?rc=250950.

Chapter 3.: Nature and Role of ESGs in Property/Casualty Insurance

Economic scenario generators have a large number of uses in property/casualty (P&C) insurance. The ability to assess values of assets and liabilities, as well as the impact of operational or strategic decisions, is predicated on being able to enumerate and describe a wide range of possible economic and financial conditions. A strategy or operational decision based on average or expected conditions may or may not look good when hypothesized against specific scenarios that deviate significantly from the expectations. In the event that such a strategy or operational decision "fails" under a particular economic scenario, the insurer will want to be aware of it, decide whether it is an acceptable risk, and consider risk management techniques to address that scenario.

We begin the chapter with an overview of certain aspects of the P/C insurance industry; the overview provides a framework for identifying some of the more important issues and economic variables that an insurer should consider when building and implementing an ESG. In the remainder of the chapter, we discuss the uses and applications of ESGs in P/C insurance, categorized into the following groups:

- Valuations of assets and liabilities
- Economic capital, regulatory requirements, and rating agency assessments
- Strategic and operational decision-making
- Risk management

This categorization is somewhat arbitrary: the uses and applications overlap significantly in this framework, and certain applications could be justified as being in more than one category.

3.1 Overview of the Property/Casualty Insurance Industry

Any technical or quantitative subject requires some understanding of context—knowledge of the domain to which one's technical skills are to be applied. Actuaries, for example, are generally good data scientists, because they often have not just technical abilities but also domain knowledge: an understanding of and appreciation for the insurance industry and the risk management environment.

In that spirit, this chapter begins with a quick review of the property/casualty insurance industry. The potential applications and usefulness of ESGs in this domain are better appreciated with a basic understanding of the insurance industry and the characteristics of its constituent companies.

3.1.1 Financial Characteristics of P/C Insurance Companies

Insurance companies tend to be relatively conservative when it comes to their investment portfolios. Much of this approach is due to regulatory requirements, which are directed toward, among other things, trying to ensure the solvency of insurers and their ability to fulfill their future obligations to policyholders. Thus, insurers are typically required to avoid certain types of risky investments in their portfolios.

Another reason for the conservative nature of insurer investment portfolios is the need for insurers to maintain a certain level of liquidity. Certain general characteristics of P&C insurers dictate many of their asset and liability cash flow patterns: the specific type of insurer, the particular lines of business the insurer writes, and so on determine the parameters and time frames associated with these cash flows. For example:

- The insurer receives premium dollars toward the beginning of the policy period—either up front or, if the policyholder pays under an installment plan, throughout the policy period.
- Certain expenses—e.g., acquisition and general expenses—are paid out by the insurer quickly from the premium proceeds. Other expenses are paid out with losses or throughout the policy period.
- The remaining premium dollars are invested.
- Accidents (or occurrences or incidents) occur during the policy period—some losses and loss adjustment expenses are paid soon after the occurrences, but some dollars may take many years to settle and be paid out.

Payments are made gradually over a period of years as losses are paid. The payment pattern can be very short (e.g., auto physical damage), longer (many liability lines), or decades (workers compensation). Patterns also may vary across insurers within a given line of business, whether due to a different mix of insureds, different claims practices, or other reasons.

It is in this environment that insurers undertake asset–liability management (ALM). By estimating when losses will be paid out—and consequently, when cash will be needed to make those payments—the insurer makes investments appropriate to that estimate. Investing premium dollars in a highly nonliquid asset, for a period of time that is much longer than the anticipated payment pattern, may necessitate liquidating certain assets and suffering significant losses as a result.

Of course, in actuality, ALM is even more complicated than that. It is one thing to treat an organization as a one-off in a simplified version of reality where insurance policies are written for one year and then that cohort of policies is followed in isolation until all losses and payments associated with those policies have been paid. But in reality, an insurer is a *going concern*, an ongoing enterprise that continues to write new business—and thus take in new premium dollars—even as it makes payments on losses associated with old policies. That continuation of business must be a part of any ALM analysis.

One direct method for managing the timings of asset cash flows so that they correspond appropriately with anticipated liability cash flows is for insurers to invest in bonds of desired effective maturity lengths, taking into account optionality. And, indeed, insurers' invested asset portfolios are generally heavily weighted toward bonds, as Tables 3.1 and 3.2 show.

Table 3.1. Invested assets (2015–2017)

Asset	P/C (%)	L/H (%)
Bonds	60.3	73.4
Stocks	23.2	2.4
Mortgage loans on real estate	0.9	11.3
Real estate	0.8	0.6
Cash/equivalents and short-term investments	6.1	2.7
Other	8.6	9.7

Note: P/C = property/casualty; L/H = life/health.

Table 3.2. Bond portfolio (2017)

Asset	P/C (%)	L/H (%)
Industrial and misc.	42.3	75.7
Special revenue	26.0	
Governments	18.3	23.2
States/territories/others	12.5	
Parent/subsidiaries/affiliate	0.4	1.1
Other	0.5	

Note: P/C = property/casualty; L/H = life/health.

While P/C insurer asset portfolios are not as heavily weighted toward bonds as life/health insurers', bonds still represented a three-fifths majority of P/C investments in 2015–2017 financial data for the industry (Table 3.1). Those bonds are distributed across all major types and issuers (Table 3.2), and, although not shown here, across maturity categories, from less than one year to more than 20 years. Common stocks, the other major component of a P/C insurer's invested assets, represent nearly one-fourth of the industry's portfolio.

Here are some other important characteristics of the P/C insurance industry:

- Insurers are leveraged, with assets generally significantly greater than surplus, so asset risk is magnified relative to the insurer's surplus or net worth.

- In accordance with generally accepted accounting principles (GAAP), assets are marked to market, but liabilities generally are not, again magnifying asset risk relative to surplus.
- Much of a portfolio's bond risk premium is for liquidity risk, not credit risk, so the spread comes cheaply if cash flow management allows bonds to be held to maturity.
- Because of the previous bullet, many bonds are held to maturity, and so asset risk is not so much characterized by trading risk but by value fluctuations.
- Policyholders tend to be risk averse and require premium discounts from insurers who carry more risk.

The ability of a P/C insurer to meet its financial obligations is a critical issue for policyholders, regulators, shareholders, other stakeholders, and society—thus, the necessity for insurers to be able to value their asset and liability portfolios regularly and effectively. Because both the major invested assets and the liabilities of a P/C insurer involve future cash flows, their valuation requires discounting with a full panel of interest rates. In addition, future liability payments are affected by future inflation of various types and degrees that may affect different sectors of the economy. Based on these financial characteristics and considerations for P/C insurers, valuing assets and liabilities requires at least the following economic factors from an ESG:

- Interest rates
    - Risk free
    - By type of bond (risk premia)
    - By maturity term (term premia)
- Inflation
    - General
    - Appropriate to line of business and type of claim payment
- Equity returns

3.1.2 Profitability: Underwriting and Operational Results of P/C Insurance Companies

Premium volumes and losses associated with many P/C lines of business are related to economic conditions, often causally. The following are some examples that have been suggested over the years:

- A relationship may exist between workers compensation losses (frequency and/or severity) and economic conditions, including unemployment. Premium levels for many lines of business—especially those with longer-tailed payout patterns—tend to decrease when interest rates increase; this is generally because of the potential for greater investment income, which leads insurers to compete on price, and may be a manifestation of an insurance pricing cycle.

- Severe catastrophic losses can have an impact on local economic conditions; combined, these effects could affect insurance capacity and prices.
- General economic conditions are generally correlated with increasing exposure levels, greater opportunities to sell policies, and written premiums. For example, if housing begins to soar due to a booming economy, exposures and premiums are likely to increase for both contractors general liability and homeowners multiple peril lines of insurance business. Strong economic activity can lead to a greater need for insurance (and insurer staff) to protect the additional wealth being produced.

Because of these and numerous other relationships between and among economic factors and underwriting/insurance variables, the ability to model a P/C insurer relative to a range of different economic conditions is critical. To the extent that an insurer is modeled as a going concern, the economic conditions generated in any given scenario produced by an ESG will lead to different underwriting and operating results—variability that the user or observer will need to take into account when assessing risks to the firm.

3.1.3 P/C Insurance Cycles

Another stylized fact about the P/C insurance industry is its tradition of exhibiting a "profitability cycle." A typical cycle might be six to 10 years long and would go from periods of high premium rates and low loss ratios (and thus relatively high profitability) to low premium rates and high loss ratios (and thus relatively low—or even negative—profitability). From an ESG standpoint, any economic or financial variable that is suspected of possibly influencing the competitiveness and profitability of the industry in a cyclical way over time should be considered for inclusion in an ESG model. (The cyclicality of P/C insurance has changed in the 21st century relative to what was seen in much of the prior century; the degree to which the cycle remains an issue is an open question.)

3.2 Overview of Property/Casualty Applications of ESGs

In the remainder of the chapter, we look at the uses and applications of ESGs to P/C insurance. As mentioned previously, we are defining an ESG model fairly narrowly—as a mechanism that simulates a variety of economic and financial variables into the future. However, we are defining the applications of ESG models fairly broadly—potentially, any application that can help model and answer questions of interest for a P/C insurer. Since many questions and decisions facing P/C insurers would benefit from an understanding of possible future economic conditions, ESGs will frequently be of value.

In the next four sections, we identify and discuss applications and uses as categorized into four groups:

- Valuations of assets and liabilities
- Economic capital, regulatory requirements, and rating agency assessments

- Strategic and operational decision-making
- Risk management

The groupings serve as an organizational principle, but this categorization is somewhat arbitrary; the various uses and applications of ESGs for P/C insurers overlap significantly, and certain applications could be justified as being in more than one category. For example, any risk management decision is also a strategic and operational decision, and calculating economic capital involves valuing assets and liabilities. Nevertheless, we feel that this scheme is intuitively appealing as the basis for discussion.

3.3 Applications of ESGs Involving the Valuations of Assets and Liabilities

3.3.1 Liabilities

The levels of losses associated with many P/C insurance lines of business are affected by general economic conditions—and sometimes by specific economic variables—in several important ways. Greater economic activity and growth can result in more exposure units and greater wealth, with the consequent need for more insurance protection. Claim severities and frequencies per unit of exposure may potentially vary as economic conditions change, and the amount and timing of loss payments are sensitive to the values of certain economic variables.

By simulating numerous future coordinated paths of economic and financial variables, ESGs can provide the basis for determining the values and sensitivities of losses to key economic and financial risks. This information can then serve as the basis for identifying and testing various financial risk management techniques.

Two economic variables in particular are worth mentioning for their potential impact on a P/C insurer's liabilities: inflation and interest rates.

- **Inflation:** A P/C insurer's liability for loss reserves can be significantly affected by the future path of inflation, as trends in cost levels over time will have an impact on the amount of losses expected to be paid. Inflationary pressures can be modeled by use of a general price change index, such as the U.S. consumer price index, or by an index whose composition reflects the specific nature of the cash flow. One extant source of such indices is the Willis Towers Watson Claim Cost Index, which is regularly published with inflation rates appropriate for a dozen or so insurance lines of business (e.g., see Pecora and Thompson 2018). This is a continuation of Masterson (1968), who created and developed this index in the 1960s. An ESG could be designed to directly model one or more of these cost indices, perhaps with an autoregressive process. Another approach would be a cascade structure, using the ESG's simulation of general inflation and then building off of those projections to yield various sub-indices.
- **Interest rates:** In addition to potential correlations with inflation, interest rates are important to the valuation of P/C liabilities, as they can be used to determine present

values of losses or other expected cash outflows. With respect to loss reserves, discounted values will still be useful information even if they are not permitted under statutory accounting—for example, for asset–liability management or possibly for GAAP accounting.

An ESG, as part of producing internally consistent future simulated paths of interest rates, inflation, and other economic variables, provides a basis for analyzing the impact of future economic volatility and measuring uncertainty. In particular, the loss reserve shown on line 1 of the liability side of an insurer's annual statement balance sheet is a single "best estimate" of future claim payments. However, that best estimate is actually a random variable that reflects different possible values with different probabilities. The loss reserve is best thought of as a stochastic value, and an actuary strives to determine (a) the best estimate, (b) the potential variability around the best estimate, and (c) how that best estimate may vary under different scenarios or conditions.

Thus, whereas it is common to think of the reserve, and indeed of insurer liabilities overall, as a point estimate, in reality it is a probability distribution with multiple levels of uncertainty: the inherent uncertainty and randomness associated with a stochastic process; institutional issues specific to the insurance and risk management industries (such as regulatory risk and social inflation); uncertainty in the emerging economic environment resulting in different expected values of future loss payments; etc.

An ESG has an important role in helping users to measure and understand the impact of this uncertainty and the potential volatility of future loss payments. This can be especially important if certain tendencies or influences underlie the loss data but have not yet manifested themselves in the data set itself.

3.3.2 Assets

In Section 3.1, we discussed the typical composition of invested assets in the P/C insurance industry. The values of nearly all such assets are sensitive—often very sensitive—to changes in economic and financial variables. Some insurers may also have assets with cash flows that are inflation dependent, such as inflation-linked bonds or inflation-correlated assets like real estate and common stocks. Even a small change in a critical economic variable can have a huge effect on the value of an insurer's asset portfolio, and thus its capital position. An ESG provides an internally consistent set of simulations that can serve as the foundation for measuring and understanding the potential volatility of asset values and cash flows.

3.3.3 Strategic Asset Allocation in an Asset–Liability Context

Some aspects of asset risk can be evaluated through a strategic asset allocation analysis. An important aspect of strategic asset allocation is developing an efficient frontier of investment classes to optimize risk and return. For example, assessing the duration behavior of the investment portfolio against the duration of liabilities on the balance sheet throughout a range of economic scenarios can lead to a deeper understanding of the effect of interest rates and other economic factors on assets, liabilities, and surplus. In theory, companies could maintain a fully liquid and essentially risk-free investment position against liability cash flow needs by investing in short-term treasuries or cash, but that is not really risk free; it is investment risk free. There is still reinvestment risk, when cash yields much less than fixed-income bonds, and business plan risk, where the strategy has a low probability of meeting the expectations of the business plan's need for investment income return. In most insurance sectors, investment income represents a significant part of overall income.

Answers to the questions of increasing return within acceptable parameters of risk may at times be found by examining alternative investment strategies that add to an insurer's portfolio what might be seen as riskier securities from a pure investment-risk perspective, but that, when considered in conjunction with an insurer's liability (underwriting) risks and ongoing cash flow needs, actually achieve a more refined efficient frontier of investment returns. When we combine the effects of asset and liability volatility by looking at change in surplus and standard deviation of surplus, we can see opportunities to reduce overall policyholder surplus risk, increase policyholder surplus, or both.

Figure 3.1. Efficient frontier comparisons

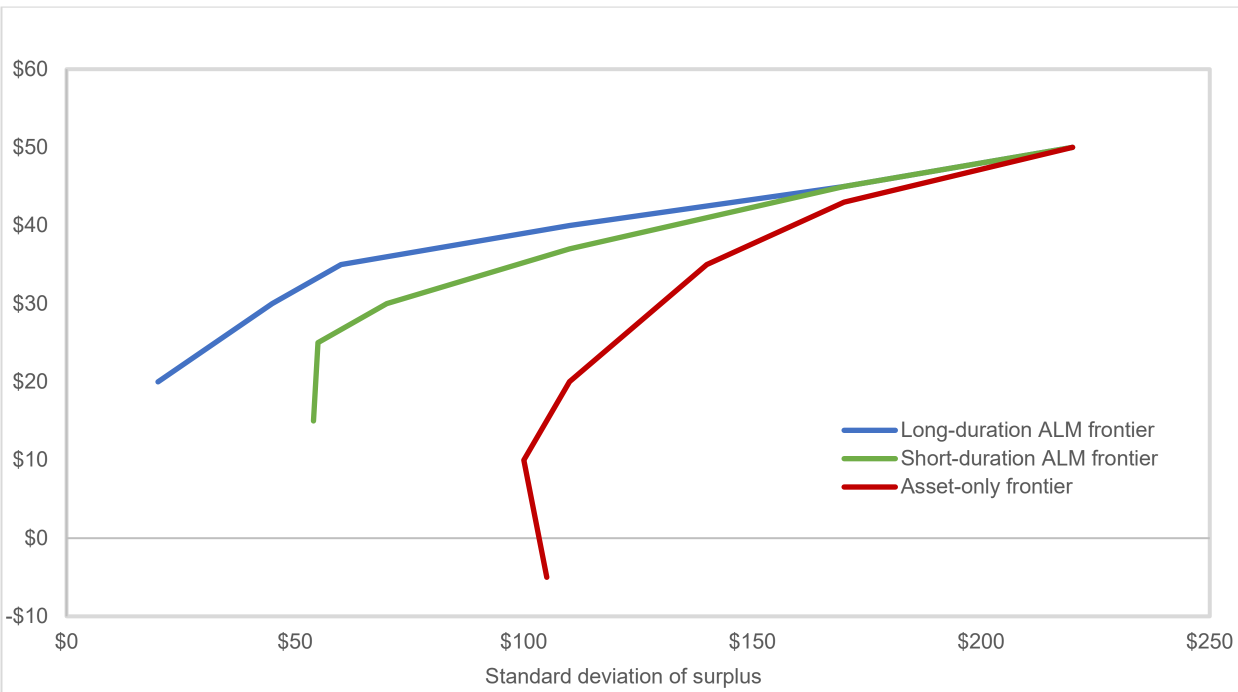

Prepared by Conning, Inc.
Note: ALM = asset–liability management.

The lowest-risk portion of an efficient frontier of an asset-only portfolio carries significant risk in an asset–liability context as liability cash flow volatility is not efficiently matched up against asset volatility. If liability cash flows are more volatile (such as might be seen in shorter-tailed automobile or property risks), the efficient frontier of investments will show greater volatility at the short or lower-risk end of the portfolio. The slope of the efficient frontier with more volatile liabilities will be steeper, as increasing asset volatility and return performance will produce marginally greater benefit in diversifying some of the liability volatility.

An efficient frontier more applicable to today's needs addresses a broader set of choices of investment allocation and also includes the recognition of a broad range of operating metric volatilities—including liability cash flows and possibly premium volatility, catastrophe risk, and even operational risk. The efficient frontier can point out potential economically efficient solutions, but this needs to be considered in the context of multiple constraints, including the need to comply with or change existing investment policy limitations, statutory accounting, and risk-based capital factors that may limit flexibility in aspiring to full "economic value" efficiency.

Figure 3.2. Strategic asset allocation efficient frontier with broader asset classes and operating volatility

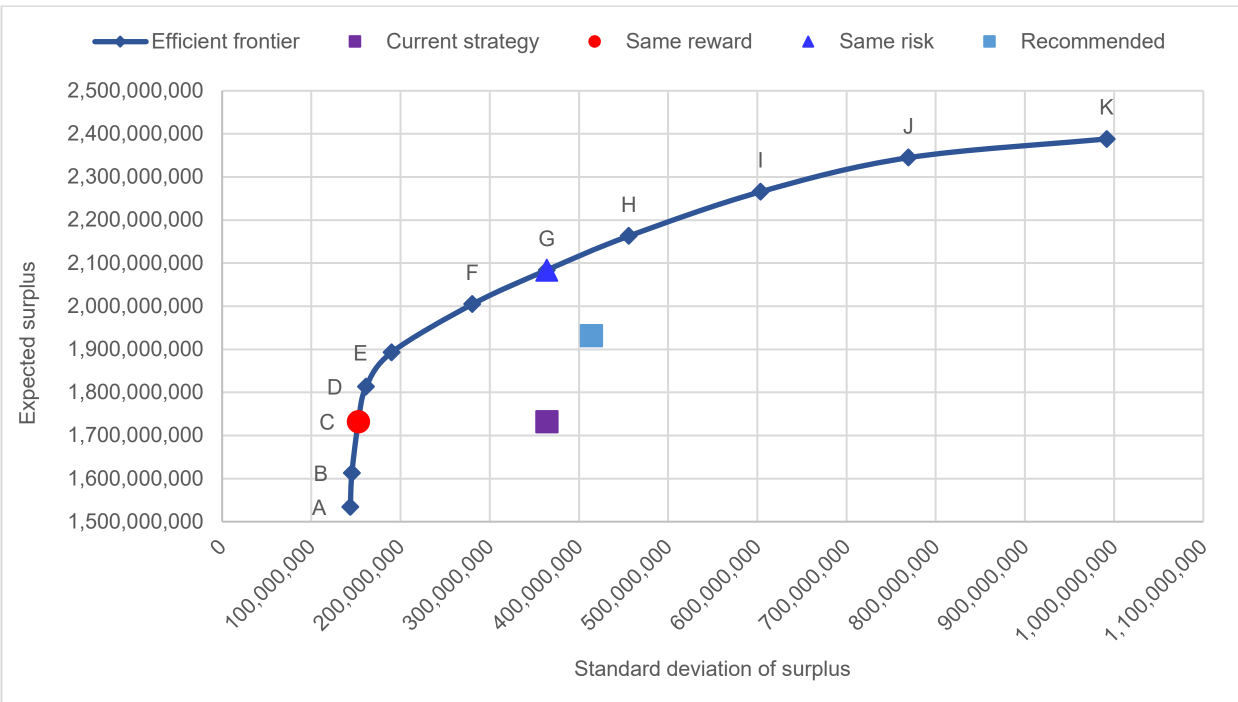

Prepared by Conning, Inc.

3.4 Applications Involving Economic Capital, Regulatory Requirements, and Rating Agency Assessments

Regulatory requirements involving assessing and monitoring an insurer's solvency largely involve the identification of situations or circumstances that can threaten the firm financially. Thus, economic capital calculations tend to be influenced by extreme, or tail, events or conditions. This is precisely the type of application at which a good ESG can excel. A company can be stress tested by identifying or flagging those paths that are deemed "extreme," and the sensitivity of the company to those types of events can be measured and examined. Having an ESG that generates internally consistent simulated future macroeconomic conditions is essential to capturing many of the circumstances that could possibly result in extreme situations.

Having a consistent set of economic conditions underlying the analysis of both assets and liabilities (or both the investment and underwriting operations) is an important aspect of economic capital modeling. Given an ESG's thorough and internally consistent simulated economic environment, both assets and liabilities can be "marked to market" as the economic capital model traces the evolution of the company over time as a going concern. This capability allows for realistic valuation projections of the company through time, providing a basis for responsible capital planning.

While adverse changes in economic and financial conditions can threaten the surplus or net worth of an insurer, for many P/C insurers the greatest perceived risk to the health of the company is often a catastrophic event, either natural or human made. Even then, however, a large catastrophe can have an impact on the general health of the economy, leading to a potential double-whammy to the P/C insurer: a large underwriting loss stemming directly from policyholder losses and a second loss due to economic conditions induced by the catastrophe. For example, inflation, an important factor for P/C insurers, could accelerate due to supply–demand issues after a major catastrophe. On the other hand, a natural catastrophe could boost certain segments of an economy—contractors and builders may see increased business after a hurricane makes landfall. Geographic areas, too, may be very differentially affected by the occurrence or nonoccurrence of a catastrophe. Capital market response with catastrophe bonds may have an impact on the cost of reinsurance and/or premium rates.

Thus, although it is generally not considered a part of an ESG, the catastrophe module used by a P/C insurer should be appropriately linked to the ESG. Subsequent economic conditions being generated after the simulation of a catastrophe should be consistent with the occurrence and level of the catastrophe. Paths with high economic expansion could also exacerbate catastrophe exposure and be relevant to reinsurance limits or other exposure considerations.

3.5 Applications Involving Strategic and Operational Decision-Making

Consider the hypothetical situation in Section 3.4, immediately above, and suppose that, for certain simulations and economic scenarios, the impact on the insurer is deemed to be "unacceptable." Perhaps, for example, the surplus falls by an unacceptable percentage in any scenario generated by the ESG in which interest rates increase by more than 300 basis points over the next five years and a hurricane strikes a heavily populated area where the company has a relatively large property insurance market share. Perhaps the need for liquidity to pay claims after a hurricane is also problematic for the company, especially if fixed-income security prices are low because of higher interest rates. Perhaps the company can handle either the interest rate increase or the hurricane—but not both.

One approach to addressing this issue might be to reduce the concentration of policies in the heavily populated area; such a strategy can be expressed in a straightforward manner—and it may well resolve the decrease-in-surplus issue—but there are some subtle operational issues involved with implementing it. Perhaps legal or regulatory restrictions exist that make it difficult for the company to reduce its insurance penetration sufficiently. There may be difficult internal battles about which policyholders to drop. Lower premiums may have suboptimal profit consequences in other economic scenarios where interest rates don't fall. Another approach might be to de-risk the future insurance exposures as well as the investment portfolio asset allocation, which could be in response to rating agency or shareholder pressures. Yet another response could be the issuance of debt or surplus notes in order to manage leverage ratios.

An ESG cannot itself make a decision about strategic or operational alternatives—but it helps provide a consistent basis for evaluating the impact of a decision across a range of different possible future circumstances. An ESG allows a company to make an informed decision.

However, to take full advantage of this modeling information, it is critical that, across an entire corporate model—for example, an enterprise risk management model—the various modules making up the full model be consistent with one another. Continuing the natural catastrophe example, an extreme or widespread catastrophic event can have effects well beyond insurance and underwriting results. Or management could be faced with very high debt issuance financing costs in periods of higher than expected inflation.

3.6 Applications Involving Risk Management

Suppose we continue the example from the prior two sections—hurricane plus interest rate increase. This time, the company is considering trying to manage the risk of both events by entering into interest rate derivative contracts to protect against the impact of higher interest rates and by purchasing additional catastrophe reinsurance for its portfolio of insurance policies in the heavily populated area.

Both of those possibilities are potential risk management solutions that a practitioner can evaluate with a model that incorporates an ESG module. Both solutions, in addition to their management of potential risks, also have costs associated with them—e.g., the ceded premium for the additional purchase of catastrophe reinsurance. To assess the relative attractiveness of the risk management solutions, the full model could be run four times: once with neither risk management solution, once with both solutions, and once each with just one of the two solutions implemented.[1] In each run, both the benefits and costs of the risk management solutions would be incorporated into the model. The full distribution of results across the same range of scenarios would allow the company to assess which of the four solution combinations is best for the company. In fact, the purchase of catastrophe reinsurance (or any reinsurance) may have an additional risk connected to the creditworthiness of the reinsurer, which may be variable with

[1] A model may be capable of running each of these solution scenarios simultaneously, by identifying and documenting multiple strategies in the outputs.

respect to significant economic stress events. Some of these stress events could correlate with significant interest rate changes.

More generally, much of current P/C insurance regulation involves assessing (and encouraging) a company's ability to manage risks, from many different sources, especially with respect to extreme (tail) events. ESGs produce future macroeconomic scenarios and joint distributions that are internally consistent and economically coherent by recognizing the interrelationships between different drivers of risks. This foundation can support a holistic approach to risk management and provides a bridge between the asset and liability (or the investment and underwriting) sides of the company. In particular, an ESG provides consistency between an insurer's investment portfolio/market risk management and its actuarial functions. In the same way that loss-reserving and ratemaking actuaries in the same organization need to communicate and ensure consistency, so too do the financial and actuarial functions need to provide a coordinated and consistent approach to the firm's risk management process. An ESG is a critical tool for facilitating this consistency.

3.7 Summary

For property/casualty insurers, the ability to assess financial statement values, as well as the impact of operational or strategic decisions, requires being able to enumerate and describe a wide range of the possible states of economic and financial conditions. Some of the more important variables that a P/C insurer should consider when building an ESG include the valuation of assets and liabilities, economic capital and regulatory requirements, strategic and operational decision-making, and risk management.

Investment portfolio decisions may be based on regulatory requirements as well as the need for maintaining a certain level of liquidity. General characteristics of P/C insurers, including prospective cash flows in the context of a going-concern enterprise, can dictate many of their asset and liability cash flow patterns, and consequently their asset–liability management decisions. Asset and liability portfolio values may be influenced by financial factors such as interest rates (risk free, risk premia, and term premia), credit risk (credit rating migration, default risk intensity), inflation (general and line-of-business specific), equity returns, and mortgage delinquency and prepayment patterns. These characteristics, along with the specific attributes and business models of individual companies, and the purpose for which the model is designed, dictate the kinds of economic and financial variables that should populate an ESG.

There are several points of intersection between P/C underwriting and operational results and the economic and financial variables generated by an ESG. For example, premium volumes and losses associated with many P/C lines of business are related to economic conditions, often causally. Furthermore, underwriting and operating factors tend to undergo significant cyclicality from periods of high premium rates and low loss ratios to low premium rates and high loss ratios. Thus, the ability to model a P/C insurer relative to a range of different economic conditions over time is critical.

Valuation of the reserves for outstanding losses (the largest liabilities of a P/C insurer) is largely the purview of actuaries. While the reserve shown on the insurer's balance sheet is a single best-estimate value, the loss reserve is actually a stochastic value with variability around the best estimate, and the best estimate may itself vary under different scenarios or conditions. A good ESG provides an actuary with a robust tool to build deeper insight into the potential volatility of future loss payments.

Other important factors in P/C balance sheet considerations include the volatility of assets (and the leverage of invested assets against surplus), the impact of foreign exchange models and multi-economy factors, and the effect of different time horizons on different line-of-business models with variable claims payout periods.

Some aspects of asset risk can be evaluated through a strategic asset allocation analysis. An important aspect of strategic asset allocation is developing an efficient frontier of investment classes to optimize risk and return. For example, assessing the duration behavior of the investment portfolio against the duration of liabilities on the balance sheet throughout a range of economic scenarios can lead to a deeper understanding of the effect of interest rates and other economic factors on assets, liabilities, and surplus.

Economic capital and regulatory requirements for P/C insurers tend to be influenced by extreme tail events, requiring responses in the form of stress testing. Often, extreme events can influence multiple aspects of the business (e.g., a catastrophic event that affects the general health of the economy), leading to a potential double impact on the P/C insurer. Inflation could also accelerate due to supply–demand issues after a major catastrophe. This is precisely the type of application at which a good ESG can excel.

Analysis of extreme events can also influence strategic and operational decision-making. An ESG cannot itself make decisions about strategic or operational alternatives, but it can provide a consistent basis for evaluating the impact of a decision across a range of different possible future circumstances.

Application of a consistent model can also provide insight into the cost and risk trade-offs of risk management questions and potential solutions.

Done correctly, an ESG can provide foundational information for making many types of corporate decisions. But to take full advantage of that modeling information, it is critical that across an entire corporate model—for example, an enterprise risk management model—the various modules that constitute the full model be consistent with one another.

References

Masterson, N. E., "Economic Factors in Liability and Property Insurance Claims Costs: 1935—1967," *Proceedings of the Casualty Actuarial Society* 55, 1968, p. 61.

Pecora, J. P., and E. M. Thompson, "2018 Claim Cost Estimates Lower than General Inflation, Reversing 2018–2017 Trend," Willis Towers Watson, https://www.willistowerswatson.com/en-US/insights/2018/11/2018-claim-cost-estimates-lower-than-general-inflation-reversing-2013-2017-trend.

Chapter 4: Perspectives on Developing and Maintaining an ESG

In this chapter we assume that the modeler has settled on the list of economic and financial variables to incorporate into the ESG, and we address the issue of how one develops the ESG and maintains its serviceability. To that end, let us think of an ESG as a collection of models used to generate a specified set of economic and financial variables. For example, a simple ESG consisting of risk-free interest rates and inflation might have a model for several maturities of risk-free interest rates and a separate model for inflation. The risk-free interest rate model would be one component model of this ESG, and the inflation model would be the other component model. Since interest rates and inflation are closely related variables, there will be some mechanism that ties the two component models together.

Developing an ESG requires that choices be made concerning the level of detail produced by each component model, the architecture of the models, and the stochastic dynamics used for each model.

Level of detail is largely governed by the use case but is a surprisingly broad area. Questions that arise about the level of detail include simulation frequency (monthly, quarterly, annually, other), maturities of simulated and reported interest rates, asset class returns, security cash flows, and default protocols for risky fixed-income instruments.

The architecture of the models involves issues such as the manner in which the models relate to one another (cascade structures, correlation assumptions, common factors), choices about whether total return is modeled directly or derived from more detailed cash flow outputs, specific institutional details of modeled variables or securities, and how auxiliary variables are computed and stored.

The selection of appropriate stochastic dynamics for a given variable is strongly influenced by the stylized facts about that variable but also involves practical trade-offs between model sophistication and the ability to efficiently maintain the model. Historical data are critical guiding factors in selecting stochastic dynamics, but a careful understanding and interpretation of the historical data also involves a measure of expert judgment.

The developer of an ESG usually addresses the issues of level of detail, model architecture, and stochastic dynamics as separate items with the recognition that there is some overlap or dependency between them.

Once the ESG has been designed and component models determined, the models must be parameterized. Parameters are usually determined relative to calibration targets and associated criteria. For market-consistent risk-neutral calibrations, parameters are usually determined by calibrating the models to market prices of financial instruments as of a certain date. Estimation of real-world models usually proceeds by determining model parameters from an analysis of historical data. The term "parameterization" encompasses both real-world estimation and risk-neutral calibration. Parameterization targets are often informed by historical data but may also be based on other business considerations, such as stress testing. Each parameterization must be validated to ensure that simulation performance is consistent with the intended model behavior and fit for its purpose.

Once an ESG has been built and incorporated into a risk management process, there is an ongoing maintenance process that involves updating initial market data and checking that the ESG continues to perform well relative to current market conditions and the intended calibration views. The following flow chart provides a general view of the entire development and maintenance process.

Figure 4.1. Process of developing and maintaining an ESG

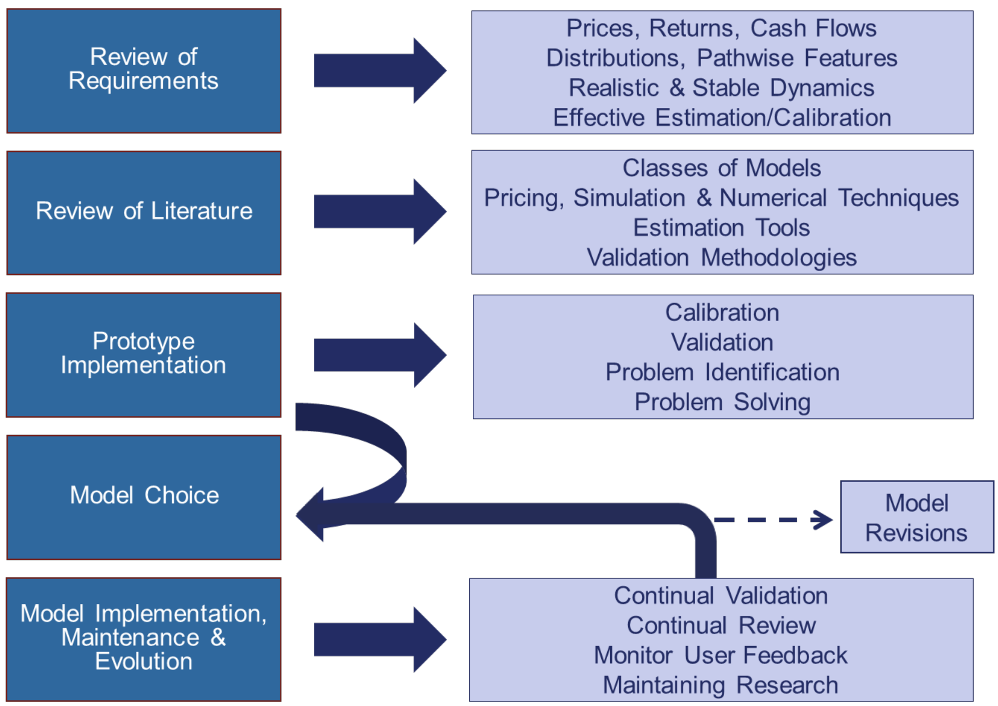

One may think about the development and maintenance of an ESG in many different ways, and different variations on the overall process can be equally successful. Regardless of the manner in which one organizes the process, the vital steps in developing and maintaining an ESG are the same. We now turn our attention to the details of those steps.

4.1 Architecture

The way in which the component models are structured and interact with one another affects the causality and correlation structure of the ESG as well as the calibration methodology. A common design approach is to use a cascade structure, an example of which is shown in Figure 4.2.

Figure 4.2. Cascade structure

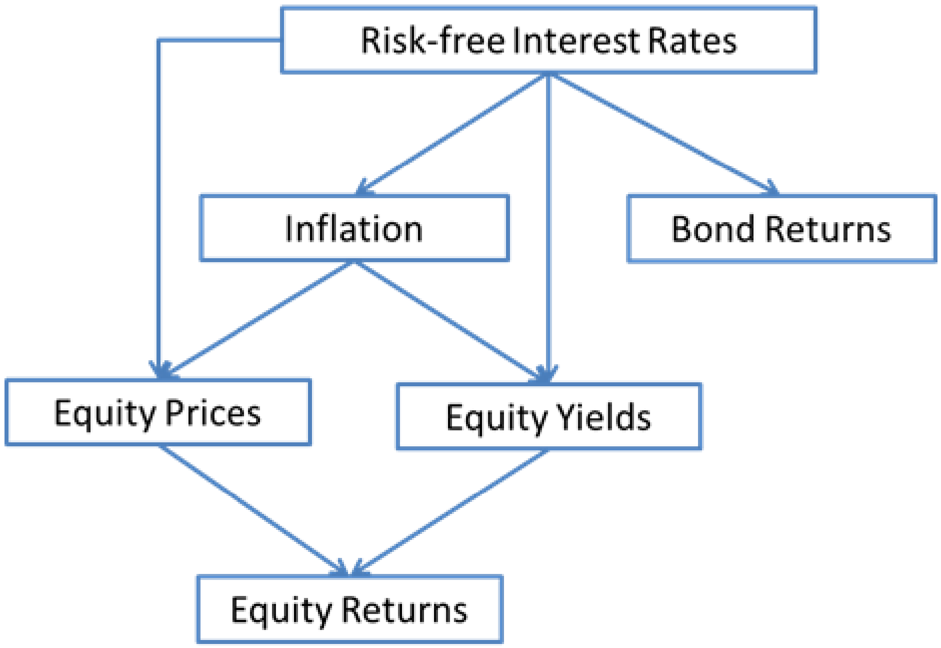

The principal idea behind the cascade structure is to provide a logical structure linking the financial and economic variables that also allows for some decoupling of dependencies so that one is not faced with the simultaneous estimation of many model parameters. In practice, an ESG can have several hundred model parameters, for which a simultaneous estimation would not be possible. The ordering of the variables is arbitrary to some extent, but risk-free interest rates are often considered as an obvious choice for the starting variable.

Some ESGs are based on econometric structures such as vector autoregressions for which it is possible to perform a simultaneous estimation of model parameters without requiring a cascade structure. An ESG based on a vector autoregressive structure can be suitable for many applications, but the limitations of such model structures are significant for some financial variables. Vector autoregressive models don't presume any theory and let the data reveal structure among variables. However, the structure implied by the data under a vector autoregressive model can be very sensitive to the selection of leg length (identification problem) and data window as well. In practical applications, vector autoregressive models tend to be used for a limited number of variables with fixed short legs. It is therefore typical that ESGs blend continuous-time and econometric models in a cascade structure.

Correlation assumptions and common factors among modeled securities and interest rates relate to the hierarchy of the causality structure of the ESG as well as calibration and parameterization methodology.

The most important outputs from an ESG are variables such as interest rates and total returns for asset classes. It is often possible to model total returns directly, bypassing the need to capture the specific cash flows for a given security type that would normally be used to compute total returns at the security level. Since some securities are complicated to model, using a total-return model can provide a useful simplification. Such a model is usually built using an available total-return index so that its historical data can be used to inform the model selection and calibration process. For example, let's assume that we are interested in modeling the returns for U.S. long-term corporate bonds. One approach is to model the multiple securities included in the asset class and calculate returns using cash flows from them. A simpler approach is to find a bond index that covers the same asset class and to build a model from the history of the index. In some applications, ESGs will use changes in interest rates and a duration assumption to mimic total returns on government bonds without specifying their cash flows. A similar approximation is sometimes used for bonds with credit risk where an additional assumption about credit costs is needed. More sophisticated investment and risk management processes often require security-level cash flows, and for those cases a total-return model would not be adequate. The decision on which asset classes can be captured with simpler total-return models is an important part of the ESG architecture.

The process of simulating ESG variables typically requires a number of intermediate steps and the computation of variables that are essential to the simulation but not normally used in the risk management process. Random draws that are used to derive underlying stochastic processes are one such example. Variables such as these are referred to as auxiliary variables. Some auxiliary variables are needed for model testing or debugging and might also be needed for model validation. The decision on which auxiliary variables are saved and available for query is also an important part of the ESG architecture.

4.2 Level of Detail

When generating ESG simulation output, one must make an assumption about simulation frequency. Does one need annual output, quarterly output, monthly output, or output at some other frequency? Continuous-time models can usually be run at any simulation frequency using the same parameters. Econometric models may need to be calibrated for the specific simulation frequency to be used. The simulation frequency options available to the user, together with the models that are available at each simulation frequency, form an important development item for an ESG.

ESG variables such as interest rates need to be computed and stored with a tenor structure. The decision as to what tenors are computed and stored is an important detail. One might decide that the tenors are set according to simulation frequency so that in an annual simulation the 1-year, 2-year, 3-year, … etc. interest rates are available. On the other hand, certain pricing applications may require a finer yield curve, perhaps with tenors incremented at three months, thereby rendering the 0.25-year, 0.50-year, 0.75-year, 1.0-year, and 1.25-year interest rates and associated discount factors available.

Some asset classes, particularly market return indices, have total returns that are naturally viewed as composed of price returns and incomes. If one wants to have this decomposition available, then the practitioner needs to develop a model for the price returns and the incomes on the securities.

Assets such as mortgage-backed securities (MBSs) and bonds with tranches (collateralized mortgage obligations, collateralized loan obligations, etc.) can have quite intricate cash flow structures from which the prices of the security must be derived. A full model for MBSs will require a prepayment function and a robust numerical procedure to compute the MBS prices from the stochastic prepayment cash flows. Tranched securities can be very hard to model because of the difficulty in modeling the factors that cause deviations from the scheduled cash flows. This may require access to a commercial vendor that has already coded the cash flow waterfall structure (typically by CUSIP number).[2]

Defaultable securities, with corporate bonds being the most common example, often have a set of assumptions governing how a default event affects the cash flows from the security, i.e., a default protocol. If one is modeling these securities, then the default protocol needs to be incorporated explicitly into the ESG, and the appropriate auxiliary variables that govern the default status of that security need to be tracked.

As a general rule of thumb, the greater the level of detail in an ESG, the slower it will run and the larger the simulation data set that is stored. One approach software vendors often use to balance the need for simulation detail versus simulation speed and efficiency is to provide some selection options allowing the user to decide which variables to store and in how much detail.

4.3 Stylized Facts and Stochastic Dynamics

To build a realistic model for a financial variable, one needs to understand the nature of the data and infer the most important properties of that financial variable from its historical behavior. At the same time, expert judgment may temper or interpret some of the attributes of the data due to market events or circumstances.

Upon examination of the historical record, one is lucky to discover a deep history for a given financial variable. When a deep history is available, one can perform a considerable amount of statistical analysis and learn a fair bit about the pathwise properties of the data-generating process. If the variable of interest has a limited amount of historical data, then one may relate it to its counterpart in another economy in order to draw inferences. This type of process involves a significant amount of expert judgment but is unavoidable when data are scarce and/or data quality is poor.

Let's look at an example where there are lots of data. In the case of U.S. short-term interest rates, about 90 years of liquid monthly data are available. The average U.S. three-month Treasury bill

[2] CUSIP stands for Committee on Uniform Securities Identification Procedures. A CUSIP number identifies most financial instruments including stocks, commercial papers, and government and municipal bonds. The CUSIP system is owned by the American Bankers Association and managed by Standard and Poor's.

yield from those data for the period January 1934 through December 2018 is about 3.48%. A chart of U.S. three-month T-bill yield shows significant variability in the data.

Figure 4.3. U.S. three-month T-bill rate, 1934–2018

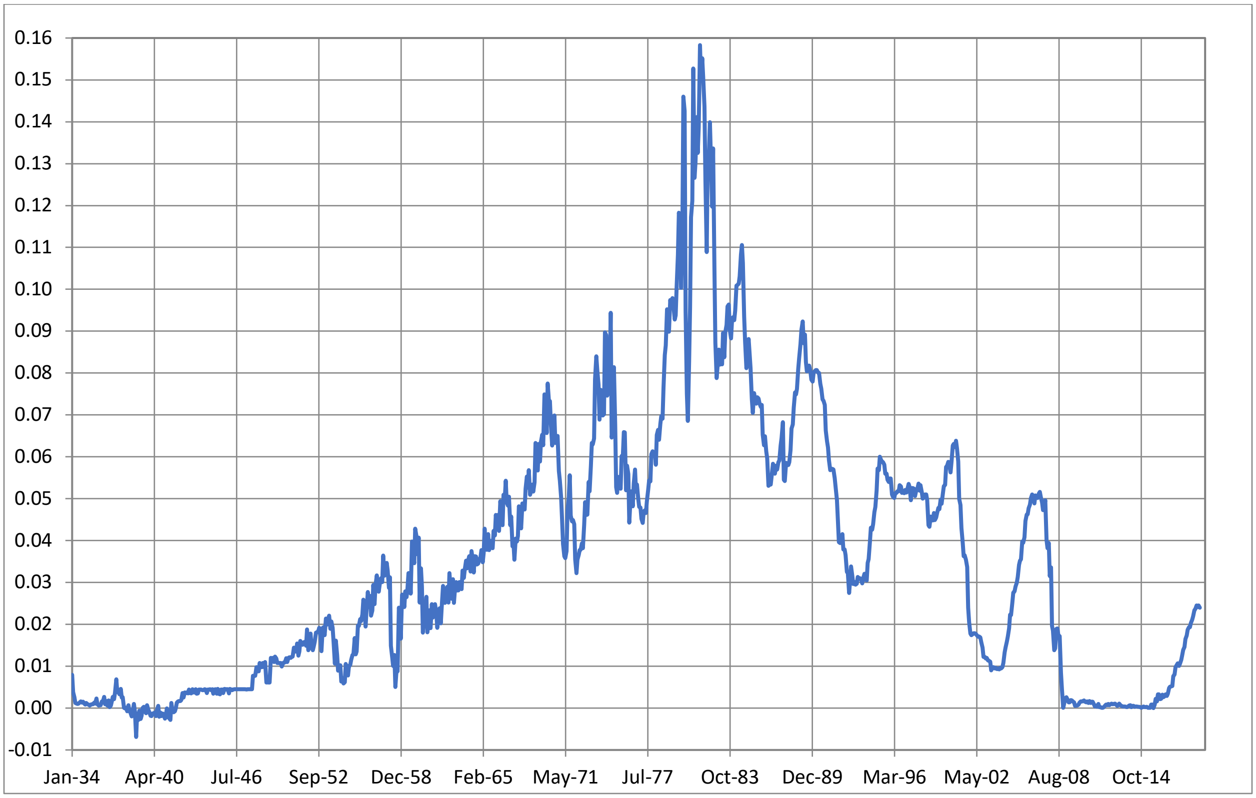

Prepared by Conning, Inc. Source: The Board of Governors Federal Reserve System, selected interest rates 1934–2019, https://www.federalreserve.gov/datadownload/Choose.aspx?rel=H15, downloaded February 2020.

One is immediately confronted with the question of whether a model for Treasury yields should allow interest rates to fluctuate over such a wide range. In particular, is the inflationary period from the late 1960s through the 1980s something that should inform the model development? The ability of a model to produce a wide range of yields while also tending to stay in a more restricted range depends both on the stochastic dynamics that are chosen as well as on the right parameters for the model. A model cannot make up for a poor choice of stochastic dynamics with a good choice of parameters.

A chart of the U.S. three-month T-bill yield for the period January 2006 through December 2018 shows a very different story.

Figure 4.4. U.S. three-month T-bill rate, 2006–2019

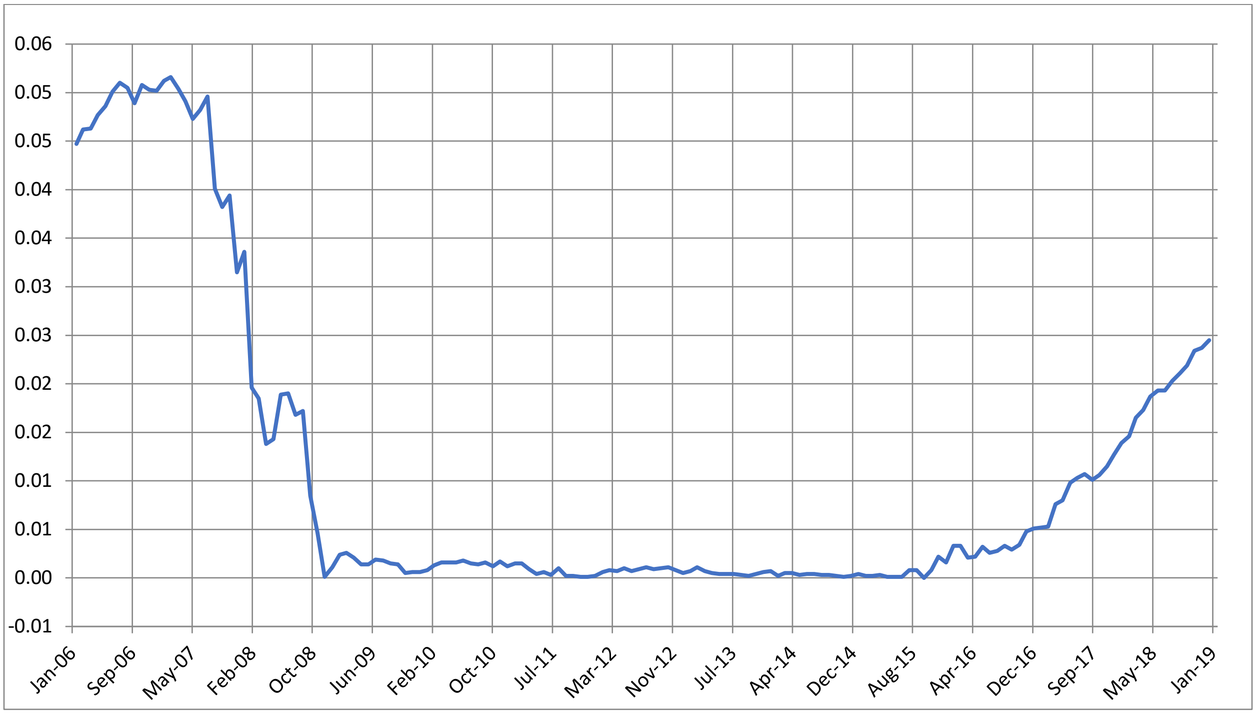

Prepared by Conning, Inc. Source: The Board of Governors of the Federal Reserve System, selected interest rates 2016–2019, https://www.federalreserve.gov/datadownload/Choose.aspx?rel=H15, downloaded February 2020.

Building an interest rate model focusing on this recent historical behavior would require incorporating some fairly sophisticated stochastic dynamics. For the interest rate to stay at a near-zero level for an extended period is relatively unprecedented. To model such behavior, new approaches such as regime-switching models or shadow-rate models have been proposed.[3] For those familiar with Japanese interest rate data, the problem is even more acute.

Another aspect of the data is their variability. The chart in Figure 4.5 shows daily changes in the U.S. three-month T-bill yield. Evidently, the recent data show some aspects of volatility clustering. Modeling approaches exist that can capture such features in the data. During the ESG design process, one needs to decide how important this feature is relative to the increased modeling sophistication needed to include it.

When we examine the variability of other interest rates, we find similar qualitative behavior, but the range and volatility depend on tenor.

[3] See Christensen (2015) for a regime-switching model that studies near-zero interest rates in recent years. Active research on shadow-rate models has been done by economists at central banks, including Christensen and Rudebusch (2015) and Bauer and Rudebusch (2015).

Figure 4.5. Daily changes in U.S. three-month T-bill rate and U.S. 10-year yield

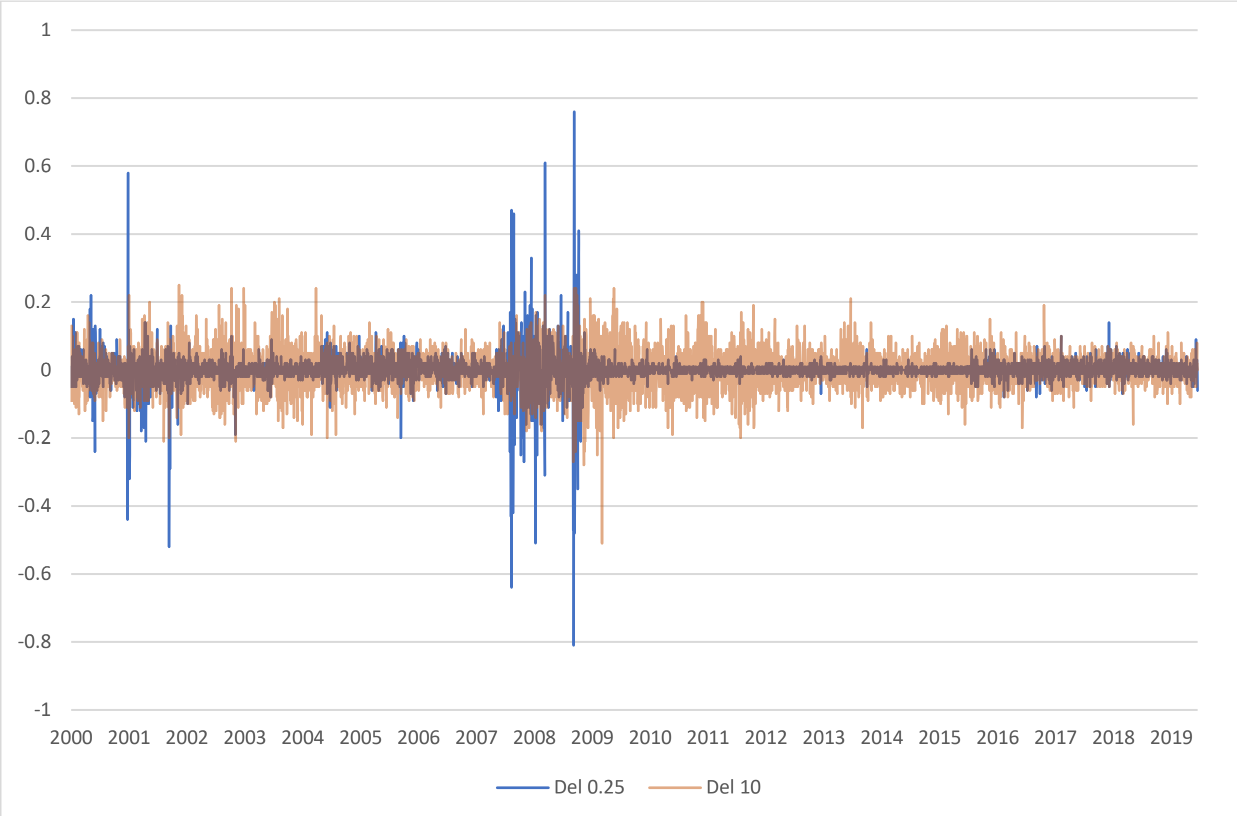

Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.

We can also look at other features of the U.S. three-month T-bill data, such as the distribution of changes in interest rates. Figure 4.6 shows a histogram of the daily change (by trading days) in the U.S. three-month T-bill rate for the period January 2000 to July 2019.

Figure 4.6. Change in three-month Treasury rate

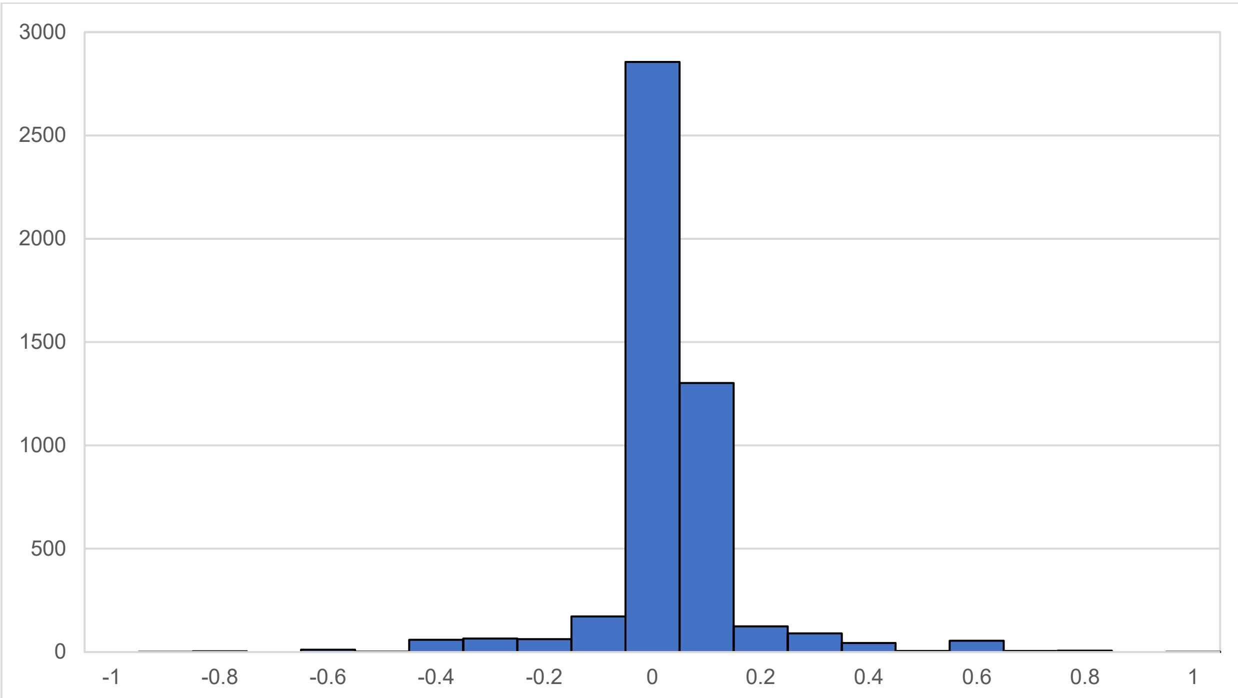

Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.

A brief examination of the properties of this distribution suggests that it is not normally distributed because there are too many observations in the tails. Capturing tails that are heavier than normal might be an important consideration for certain use cases; on the other hand, one does not want the model to generate a preponderance of interest rate paths that wander to extremes. Understanding the key features of the data-generating process is one thing; finding the right stochastic dynamics to craft a useful and stable model can be another.

For some asset classes, stochastic volatility or jumps in price are a defining feature. Figure 4.7 shows historical daily returns for the S&P 500; it is clear from this chart that the volatility of the daily returns is not constant over time and there are occasional changes that may be too large to be parts of a continuous process. If a model for these asset classes does not incorporate stochastic volatility or jumps in price, then the model would produce unrealistic outputs. Equity indices are the classic example of this phenomenon.

Figure 4.7. S&P 500 daily returns

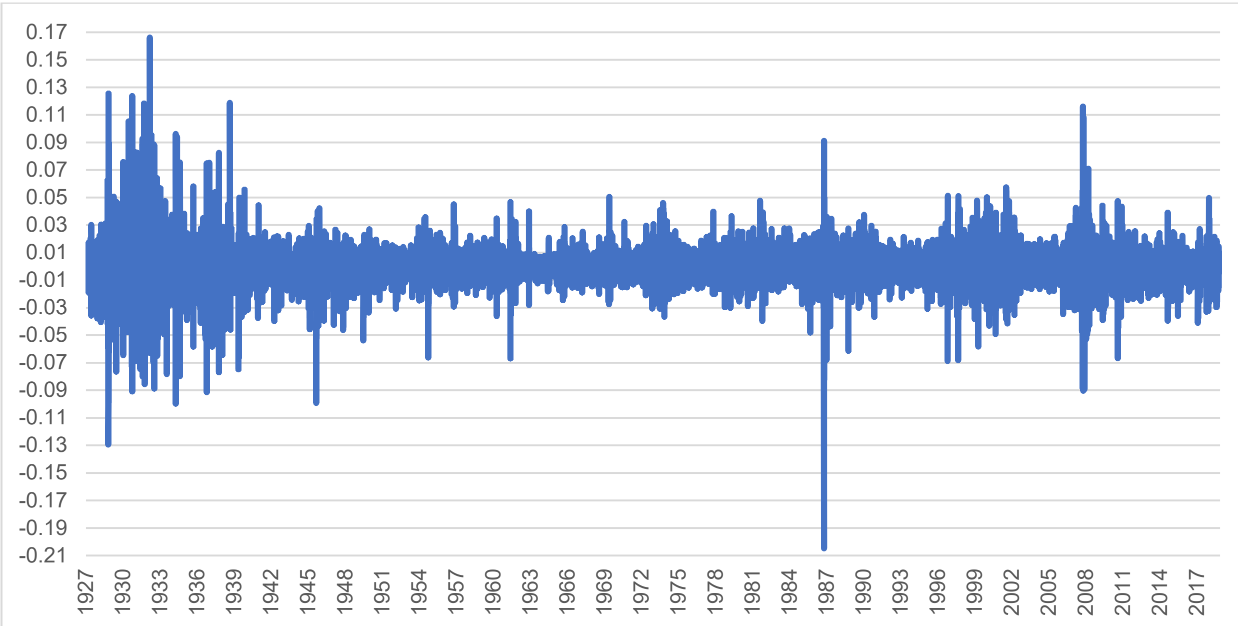

Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.

The standard geometric Brownian motion (GBM) model, which underpins the Black–Scholes (1973) model for equity prices, assumes that equity returns are log-normally distributed. This leads to simulated daily returns as shown in Figure 4.8, and it is clear that the stochastic dynamics of this model are not sufficient for the job.

Figure 4.8. Simulated returns with geometric Brownian motion model

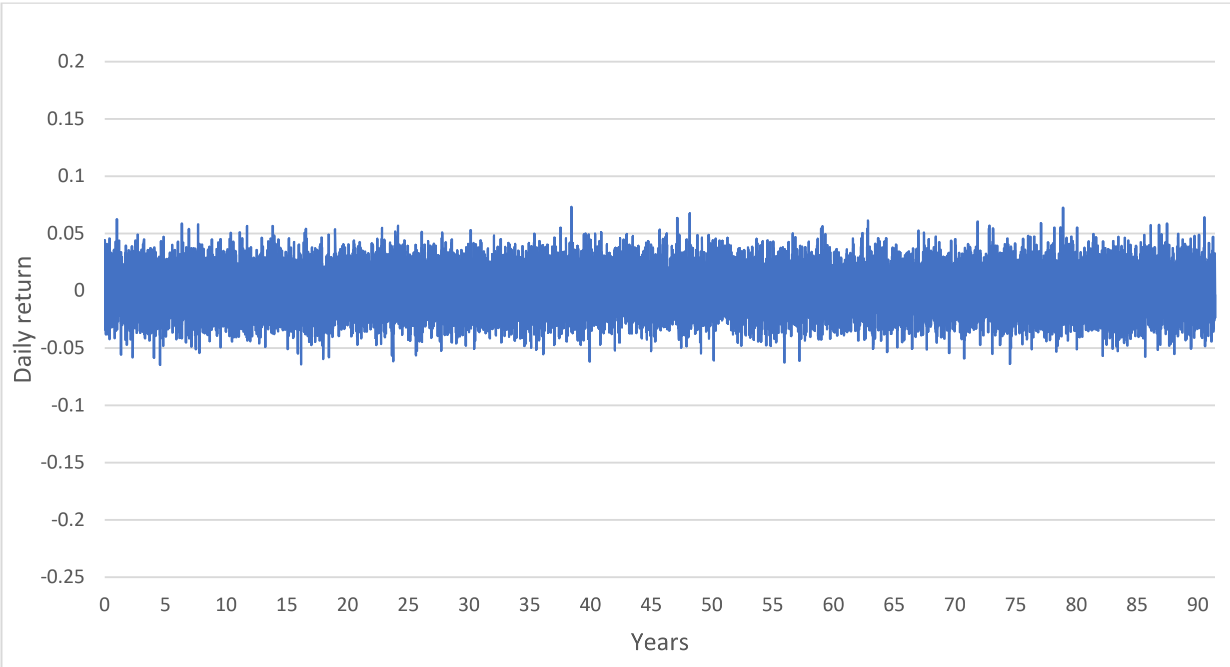

Prepared by Conning, Inc.

If stochastic volatility is added to the model (but not jumps), the model will produce simulated behavior such as we see in Figure 4.9. That chart shows a scenario of stock returns generated by Heston's stochastic volatility model (see Chapter 11 for an explanation of Heston's paper).

Figure 4.9. Simulated returns with stochastic volatility model

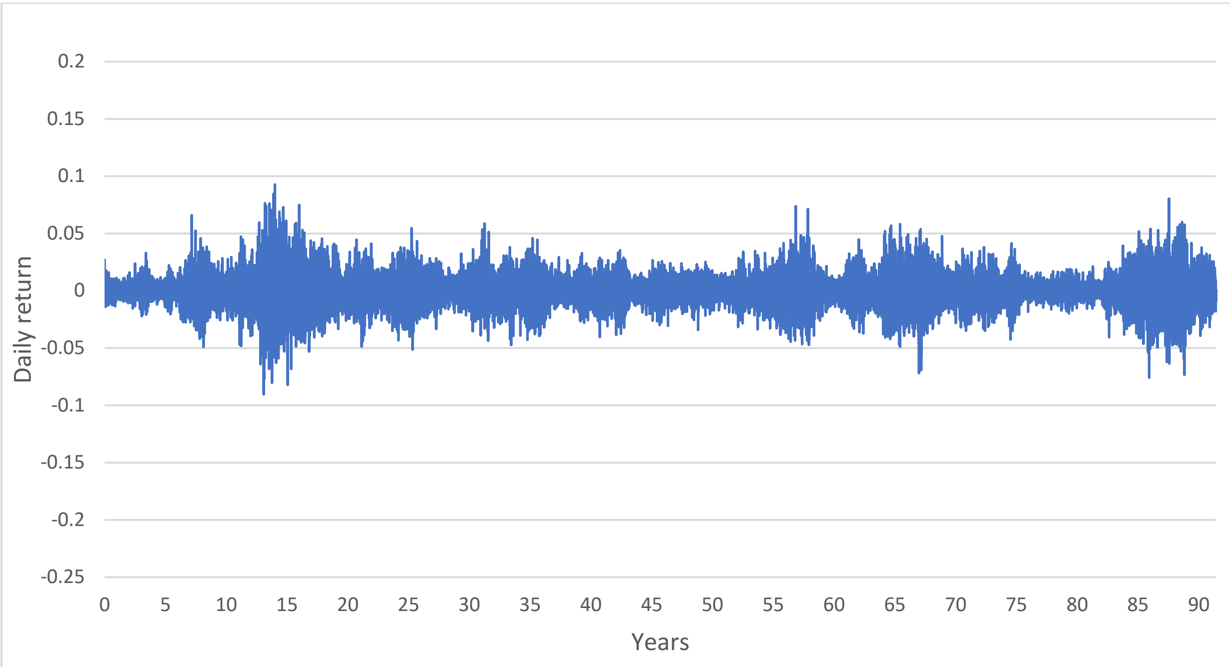

Prepared by Conning, Inc.

The Heston model generates more extreme events than the GBM model, and this is closer to what is seen in the data. It also produces epochs of high and low volatility. The Heston model captures the volatility clustering and the correlation between price and volatility. One reason the Heston model is so popular among the class of stochastic volatility models is that it provides a closed-form solution for option pricing when the asset price is correlated with stochastic volatility. The Heston model is capable of producing a variety of volatility smiles (U-shaped pattern of implied volatilities across different strikes) and volatility term structures (convex-shape implied-volatility curve across option maturities).

But in order to capture the same overall volatility as seen in the historical data, the stochastic volatility model without jumps tends to produce the extreme events too frequently. Only when a jump component is added can the contribution of the stochastic volatility process be muted to mimic the overall volatility structure of the S&P 500 returns. By using stochastic dynamics that include both stochastic volatility and jumps, it is possible to obtain the simulated behavior shown in Figure 4.10. The scenario shown there is generated with a "stochastic volatility with jump" model suggested by Bates (1996, see Chapter 11 for an explanation).

It is often observed that the price of options that protect against large price drops (e.g., out-of-the-money put option) becomes unusually high. This phenomenon is often referred to as a *volatility smirk,* and it is related to the fear of a market crash. The Heston model will capture the high volatility (i.e., high option price) by increasing the overall volatility level, which results in generating extreme events too often. The stochastic volatility with jump model is more flexible and can lead to more accurate option pricing. By letting the jump part explain the big changes, the continuous diffusion part can generate a more realistic result.

Figure 4.10. Simulated returns with stochastic volatility with jump model

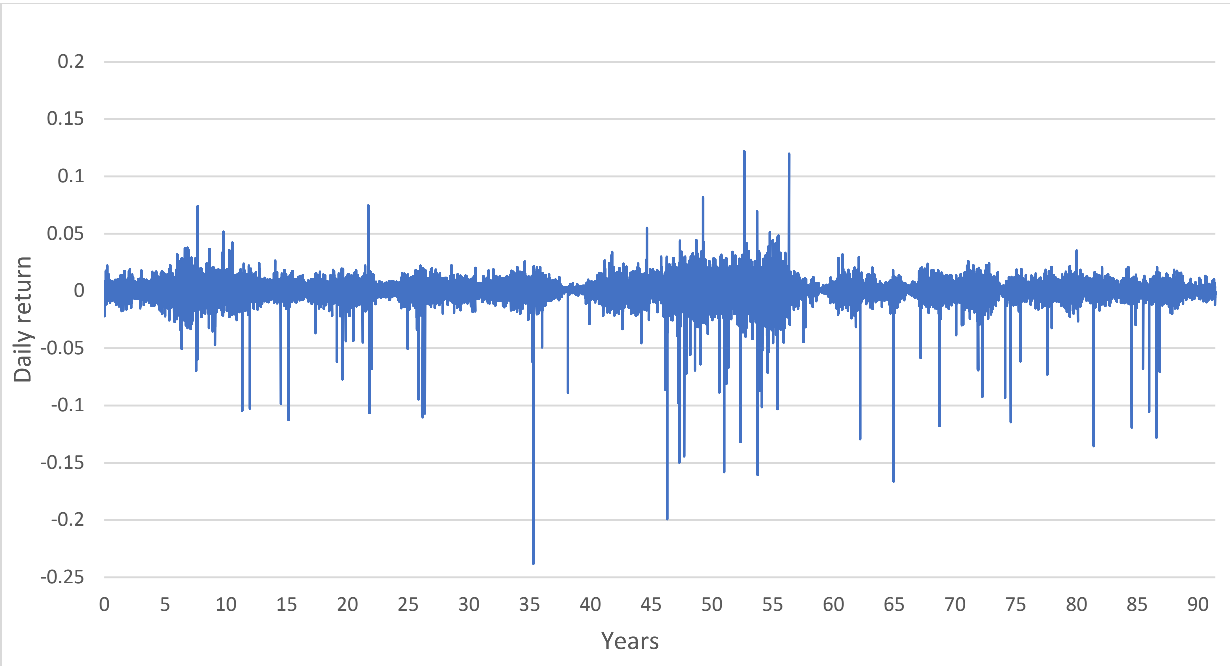

Prepared by Conning, Inc.

When developing an ESG, a comprehensive understanding of the data-generating processes for the financial variables to be modeled is recommended. Finding a stochastic dynamic that is suitable for modeling a given financial variable is rarely easy; even in cases where sophisticated stochastic models are available, challenges in calibrating or simulating such a model may result in the selection of a lesser but adequate model.

4.4 Data Sources

The availability of data sources is critical for developing and maintaining an ESG. Historical data are the standard for setting stylized facts and targets that can be used for ESG model calibration and validation.

Some financial data are available for free, particularly risk-free interest rate data and macroeconomic data such as inflation, unemployment rate, and GDP data. The Board of Governors of the Federal Reserve System, for example, provides a host of interest rate data series to the general public (https://www.federalreserve.gov/releases/h15/). Among them, nominal (non-inflation-indexed) yields on Treasury securities at "constant maturities," or constant maturity Treasury yields (CMTs), are probably the most important.[4]

The zero-coupon yield is the yield to maturity on a non-coupon-paying bond (zero-coupon bond). The zero-coupon yield curve is critically important in security pricing but is not directly observable. In the case of the United States, there are no tradable zero-coupon long-term Treasury securities. U.S. savings bonds are zero-coupon bonds issued by the U.S. Treasury but are not tradable. A zero-coupon yield curve can be estimated from a coupon curve using a recursive bootstrapping method.

Inflation data and unemployment rates are published monthly by the U.S. Bureau of Labor Statistics (https://www.bls.gov/). GDP and other related statistics are published quarterly by the Bureau of Economic Analysis (https://www.bea.gov/).

Data such as corporate bond yield and return will most likely require a paid subscription to a data vendor.

The frequency at which data are available can vary. Financial data, such as interest rates or return series, are generally available at a daily frequency for recent time periods. GDP data are released at a quarterly frequency while consumer price index and the unemployment rate data are available at a monthly frequency. Less current financial data as well as data for some alternative asset classes might be available only quarterly or annually. Data that are available at a higher frequency are usually desirable for ESG development and maintenance. It is typical to use data at monthly or quarterly frequency, but some applications call for daily frequency, particularly when analyzing the volatility of jump behavior.

For an ESG that is part of an ongoing risk management program, it is essential that the data sources are accurate, regularly updated, and capable of automated download. This is because, in this use context, the ESG is likely to be recalibrated, validated, and initialized to market conditions at regular intervals, such as year-end or quarter-end. It is also important for ESG users to refer to common sources of data. Most data vendors use a ticker system that identifies the data series. Anyone who knows the ticker for a given financial variable can be sure they are getting the same data as others using that ticker.

Commercial data vendors that offer subscription services for a broad range of financial and macroeconomic data include

[4] These yields are calculated by the U.S. Treasury from market yields (more precisely, closing market bid yields) on actively traded (on-the-run) securities. Currently the constant maturities are 1, 3, and 6 months and 1, 2, 3, 4, 5, 7, 10, 20, and 30 years. Constant maturity yields of inflation-indexed securities (TIPs) are also available from the same website for constant maturities of 5, 7, 10, 20, and 30 years. CMT yields of maturities beyond one year are yields on coupon-paying Treasury notes and bonds and those of maturities less than one year are yields on non-coupon-paying Treasury bills.

- Bloomberg,
- Thomson Reuters, and
- Global Financial Data.

Some of these vendors provide automated electronic data delivery services as well.

Barclays Capital Live has a comprehensive list of fixed-income yields and returns.

Thomson Reuters delivers a comprehensive list of muni bond data through its Municipal Market Monitor portal.

Duff and Phelps annually publishes the *SBBI Yearbook*, formerly published by Morningstar.

Many central banks and other government organizations provide economic data. The following are examples of such data sources:

- St. Louis Fed (FRED) Database
- Board of Governors of the Federal Reserve System
- Bank of Canada
- Statistics Canada (CANSIM)
- Bank of England
- Ministry of Finance (Japan)
- Swiss National Bank
- German Bundesbank
- U.S. Bureau of Labor Statistics
- U.S. Bureau of Economic Analysis

There is a considerable amount of useful data scattered in the academic literature. It may appear in journal articles, books, or websites.

Bond-rating services, such as Moody's, Standard & Poor's, Fitch, and DBRS, provide periodic reports containing data useful for modeling corporate bonds. Such data include rating transitions, default probabilities, recovery rates, and combined credit costs.

Here are a few other miscellaneous data sources:

- MeasuringWorth.com (https://www.measuringworth.com/)
- Robert Shiller's data page (http://www.econ.yale.edu/~shiller/data.htm)
- Elroy Dimson, Paul Marsh, and Mike Staunton, *Triumph of the Optimists: 101 Years of Global Investment Returns,* Princeton, NJ: Princeton University Press, 2002

- Canadian Institute of Actuaries, "Report on Canadian Economic Statistics"
- Øyvind, Eitrheim, Jan T. Klovland, and Jan F. Qvigstad, eds., "Historical Monetary Statistics for Norway 1819–2003," Norges Bank Occasional Papers No. 35
- Øyvind, Eitrheim, Jan T. Klovland, and Jan F. Qvigstad, eds., "Historical Monetary Statistics for Norway—Part II," Norges Bank Occasional Papers No. 38

4.5 Parameterization/Calibration Process and Methodology

Once the ESG component models have been selected, the next step is the parameterization/calibration process of selecting model parameters. Estimating model parameters with historical data and calibrating models to specific market conditions are key parts of the process. If the component models have been designed well, there should be model parameters that allow the models to perform as required.

An important aspect of the calibration process is setting (long-term) targets for modeled variables. The choice of calibration targets unavoidably reflects a view of longer-term trends of the variables. Views can be of an explicit nature, such as the specification of average interest rate levels at various simulation horizons, or views can be of an implicit nature wherein a long-run interest rate target might be adopted from historical data; in the latter case, the behavior of the model at shorter simulation horizons is determined by the rate of mean reversion of the model.

Different model calibration may be required for different applications. For example:

- **Normal-course business management:** Management adopts what it deems to be a reasonable set of assumptions for investment returns, inflation, and risk. Some adjustments may be made to the calibration to capture a worst-case or best-case view, and some additional simulations may be run to understand sensitivities across a plausible set of assumptions.

- **Regulatory requirements:** Regulators may require models to comply with specified calibration criteria, such as percentiles or tail behavior. Certain regulatory regimes—such as Article 226 of the European Insurance and Occupational Pensions Authority (EIOPA) Solvency II Directive—mandate that firms take into account the assumptions of their internal model in making business decisions. Clarification of the underlying assumptions of the economic modeling is a key element in satisfying such regulatory requirements, with calibration properties being a key issue.

- **Stress testing:** The process of stress testing has been widely adopted in risk management. Stressing the core assumptions of a model represents an alternative view that allows for the impact of extreme and unexpected events to be studied, quantified, and potentially mitigated.

- **Investment management:** ESGs are a very useful tool for investment management whether one applies a basic risk–reward analysis or a sophisticated holistic approach such as strategic asset allocation (SAA). All approaches are based on the analysis of expected returns and risk from an ESG. The SAA calibration may embed the shorter-term assumptions of the internal model into a multiyear stochastic projection, and/or it may reflect a long-term steady-state view of the economy and financial markets.

The robustness of the scenarios depends on achieving parameterization objectives, obtaining a range of scenarios that encapsulate historical experience, and generating some extreme but plausible scenarios.

The general process for parameterization and calibration requires a validation check to ascertain that the criteria are met (Figure 4.11).

Figure 4.11. Real-world parameterization process

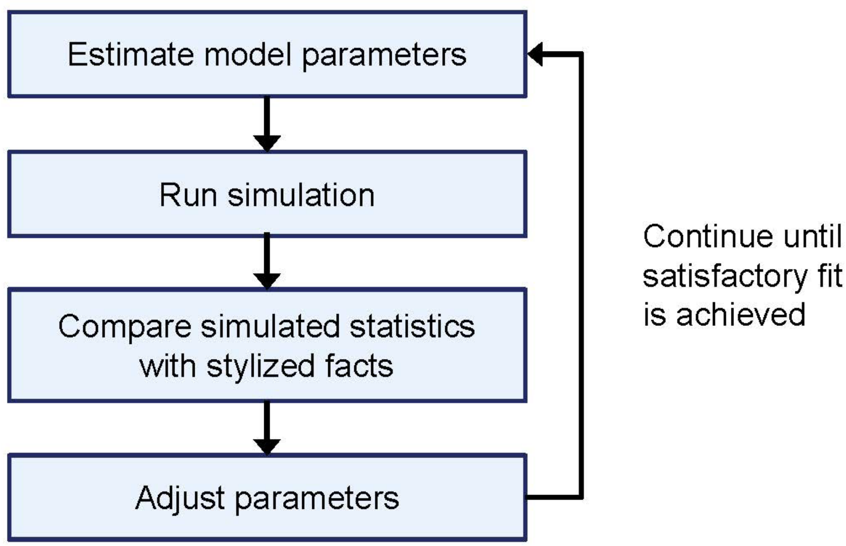

4.5.1 Model Parameter Estimation

This subsection discusses briefly the most common methods of model parameter estimation.

Estimation of econometric models such as vector autoregressive models is a straightforward extension of estimation of univariate autoregressive models, and, with some modifications, one can use most of the analytic tools from univariate model estimation. Under standard model specifications, the model parameters can be estimated by ordinary least squares. If the modeled variables are assumed to be normal, the least square estimator is equivalent to the maximum likelihood estimator. For more general model specifications, such as non-normal error distribution, maximum likelihood estimation (MLE) is a standard approach. When the distribution of modeled variables is not known, quasi–maximum likelihood estimation is commonly applied by assuming that the underlying distribution is normal. A good reference for econometric analysis is Hamilton (1994). Estimation of continuous-time models is more challenging, mainly because in many cases one does not know the exact distribution of discretely observed samples. Some diffusion models, such as the GBM model, the Ornstein–Uhlenbeck process, and the Cox–Ingersoll–Ross (CIR) model, are exceptions for which the transition densities (conditional distributions) of observed samples are analytically available and admit MLE.

Another layer of difficulty stems from the fact that some variables are unobservable. The volatility process in the Heston model or the CIR factors in the multifactor interest rate model are examples of unobservable processes. MLE cannot be applied directly to parameter estimation of those models even if the transition densities of underlying processes are known. Kalman filtering (more precisely, estimation based on a Kalman filter) is a method one can use for parameter estimation when state variables are unobservable. Detailed discussion of Kalman filtering is beyond the scope of this document but the key idea is as follows. A filter can be seen as a mechanism that produces estimates of the hidden states of a system using all information available from observations (and control variables if available). Among many filters, the Kalman filter is the best linear filter in the sense that it minimizes the mean squared error of the estimates of the unobservable state variables. The estimates of the state variables, or filtered states, are now fed into the system to generate forecasts of the observed variable. Parameter estimates are calculated by maximizing the likelihood of the difference (the measurement error) between the observed values and the forecast. Typically, the measurement errors are assumed to follow normal distribution. Chen and Scott (2003) discuss in detail the estimation of two-factor and three-factor interest rate models using a Kalman filter.

When the likelihood function is not available explicitly, less efficient estimators can be obtained by optimizing other criteria. For example, the generalized method of moments (GMM) utilizes certain moment conditions rather than the full density and minimizes some distance between theoretical moments (as functions of parameters) and their empirical counterparts. Suppose that a stationary process y_t follows the stochastic differential equation

$$dy_t = \alpha(y_t;\theta) + \sigma(y_t;\theta)dW_t$$

and parameter vector θ satisfies the following moment conditions,

$$E\left(g_i(\theta; y_t, y_{t-1}, \cdots, y_{t-p})\right) = 0,$$

for $i = 1, 2, \ldots$, k, where E denotes mathematical expectation. For example, if y_t is a process with $E(y_t + y_{t-1}) = \theta_1 + \theta_2$, then

$$g_i(\theta; y_t, y_{t-1}, \cdots, y_{t-p}) = (\theta_1 + \theta_2) - (y_t + y_{t-1}).$$

If the number of moment conditions k is equal to the size of parameter vector θ, the system of equations can be solved for θ, and the classical method of moments can be used. GMM is developed to utilize the moment conditions efficiently when there are more moment conditions than the number of parameters.

The idea behind GMM is to use a sample counterpart of the above moment conditions,

$$\hat{g}_i(\theta,; y_t, y_{t-1}, \cdots, y_{t-p}) = \frac{1}{T}\sum_{t=1}^{T} g_i(\theta; y_t, y_{t-1}, \cdots, y_{t-p}),$$

and to choose θ so as to make the sample moment as close as possible to zero:

$$\hat{\theta}_{GMM} = Minimize_{\theta}\left\{\hat{g}\left(\theta,; y_t, y_{t-1}, \cdots, y_{t-p}\right)' \Omega \hat{g}\left(\theta,; y_t, y_{t-1}, \cdots, y_{t-p}\right)\right\},$$

where Ω is a matrix weighing importance of moment conditions. For a general discussion of GMM, see Hamilton (1994).

GMM requires that the theoretical moments are known explicitly as functions of parameters. That is, the moment conditions can be written as

$$g_i\left(\theta; y_t, y_{t-1}, \cdots, y_{t-p}\right) = h_i\left(y_t, y_{t-1}, \cdots, y_{t-p}\right) - m_i(\theta),$$

where $m_i(\theta)$ is known and suitable for optimization. The idea of the simulated method of moments (SMM) is that when the function $m_i(\theta)$ is unknown, an approximation from simulation can replace it. Simulated moments can be calculated from simulated values of the process y_t^S using the stochastic differential equation

$$m_i^S(\theta) = \frac{1}{L}\sum_{t=1}^{L} g_i\left(\theta; y_t^S, y_{t-1}^S, \cdots, y_{t-p}^S\right).$$

The simulated sample size L consists of multiples of the observation size T. SMM estimates can be formulated as

$$\hat{\theta}_{SMM} = Minimize_{\theta}\left\{\hat{g}^s\left(\theta,; y_t, y_{t-1}, \cdots, y_{t-p}\right)' \Omega \hat{g}^s\left(\theta,; y_t, y_{t-1}, \cdots, y_{t-p}\right)\right\},$$

where

$\hat{g}^s\left(\theta,; y_t, y_{t-1}, \cdots, y_{t-p}\right) = \frac{1}{T}\sum_{t=1}^{T} h_i\left(y_t, y_{t-1}, \cdots, y_{t-p}\right) - m_i^S(\theta)$. SMM requires little analytical tractability but instead depends on numerical simulation of the model and is computationally intensive. Gouriéroux and Monfort (1997) discuss SMM and related methods extensively.

Another simulation-based method is the Markov chain Monte Carlo (MCMC) approach combined with Bayesian statistical inference. Let's consider the GBM model as an illustrative example. Suppose that an asset price follows a GBM:

$$dS_t = \left(\mu + \frac{1}{2}\sigma^2\right) S_t\, dt + \sigma\, S_t\, dW_t,$$

where μ is expected return and σ is the volatility. This model has a closed-form solution in return form

$$z_t \coloneqq \log\left(\frac{S_t}{S_{t-1}}\right) = \mu + \sigma\varepsilon_t,$$

where $\varepsilon_t \sim N(0,1)$. For a given sample $Z = (z_1, z_2, \cdots, z_T)$ of returns, the (conditional) likelihood is given by

$$\mathrm{P}(Z|\mu,\sigma^2) = \left(\frac{1}{\sqrt{2\pi\sigma^2}}\right)^T exp\left(-\frac{1}{2\sigma^2}\sum_{t=1}^{T}(z_t-\mu)^2\right),$$

and the MLE is attainable and is considered as the best estimator for the parameters (μ, σ^2) in the classical school of statistics. The Bayesian school takes a different approach, and this subsection provides a brief explanation of the Bayesian method and the role MCMC plays.

Whereas the classical inference focuses on finding point estimates of parameters and their statistical characteristics, the goal of the Bayesian method is to find the posterior distribution of parameters by combining information from observations and prior distribution of parameters. If the (posterior) distribution $P(\mu, \sigma^2|Z)$ of the GBM parameters (μ, σ^2) is available, the object of interest may be

$$E[f(\mu,\sigma^2)|Z] = \int f(\mu,\sigma^2)\ P(\mu,\sigma^2|Z)d(\mu,\sigma^2)$$

for some function $f(\cdot)$. Simple summary statistics such as posterior mean, posterior variance, and posterior quantiles fall in this category, but the function can be more complex. The idea of MCMC estimation is to get samples of parameters $\left\{\mu^{(g)}, \sigma^{2,(g)}\right\}_{g=1}^{G}$ from the posterior distribution to calculate

$$E[f(\widehat{\mu,\sigma^2})|Z] = \frac{1}{G}\sum_{g=1}^{T} f\left(\mu^{(g)},\sigma^{2,(g)}\right).$$

The Monte Carlo samples from the posterior distribution are not independent but form Markov chains, and this guarantees that the distribution of samples converges to the true posterior distribution as $G \to \infty$. Assuming that the prior distributions of parameters $P(\mu)$ and $P(\sigma^2)$ are given, the Bayes rule implies that

$$\mathrm{P}(\mu|\sigma^2,Z) \propto P(Z|\mu,\sigma^2)P(\mu),$$

$$\mathrm{P}(\sigma^2|\mu,Z) \propto P(Z|\mu,\sigma^2)P(\sigma^2).$$

For some choices of priors, these posterior distributions admit inversion and sampling. For example, if the prior $P(\mu)$ is normal and $P(\sigma^2)$ is inverse gamma, then it is known that the posteriors are also normal and inverse gamma. The priors with this property are called conjugate-priors. The MCMC sample for the parameter $\left\{\mu^{(g)}, \sigma^{2,(g)}\right\})_{g=1}^{G}\}$ is drawn from the marginal posterior by

$$\mu^{(g+1)} \sim \mathrm{P}\left(\mu \middle| \sigma^{2,(g)}, Z\right),$$

$$\sigma^{2,(g+1)} \sim P\left(\sigma^2 \middle| \mu^{(g+1)}, Z\right).$$

For complex models, the calculation of posterior distributions and their inversion can be intractably difficult, and the use of conjugate-priors for computational reasons can limit the use of MCMC methods. Many approaches have been proposed regarding sampling methods and selection of priors. Robert and Casella (2004) provide a good textbook treatment of Bayesian inference and MCMC methods. Johannes and Polson (2010) discuss MCMC estimations for many continuous-time financial models.

4.6 Validation Process

The validation process involves checking that the calibrated model performs in line with the calibration criteria and that the general behavior of the model is consistent with the stylized facts. Checking that the model is performing in line with the calibration criteria will usually include a comparison of the simulated model statistics with the calibration targets. One way to make such a comparison is with a fan chart such as is shown in Figure 4.12.

Figure 4.12. Fan chart—10-year Treasury

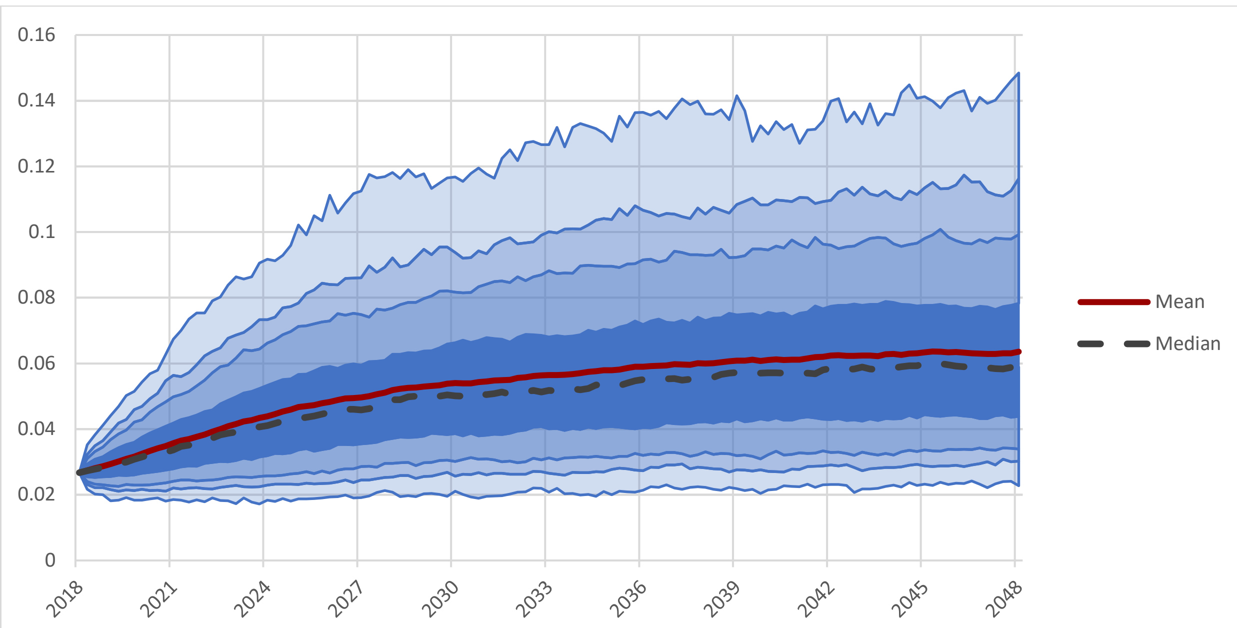

Percentiles plotted: 1%, 5%, 10%, 25%, 50%, 75%, 90%, 95%, and 99%. Prepared by Conning, Inc. Source: GEMS® Economic Scenario Generator scenario.

In such an analysis, one can clearly identify the targets and directly compare the average simulated values with those targets. At the same time, one can immediately obtain a range for the simulated variable and percentiles at various simulation horizons.

A related type of chart such as that shown in Figure 4.13 is also a useful validation aid.

Figure 4.13. Histogram—10-year Treasury

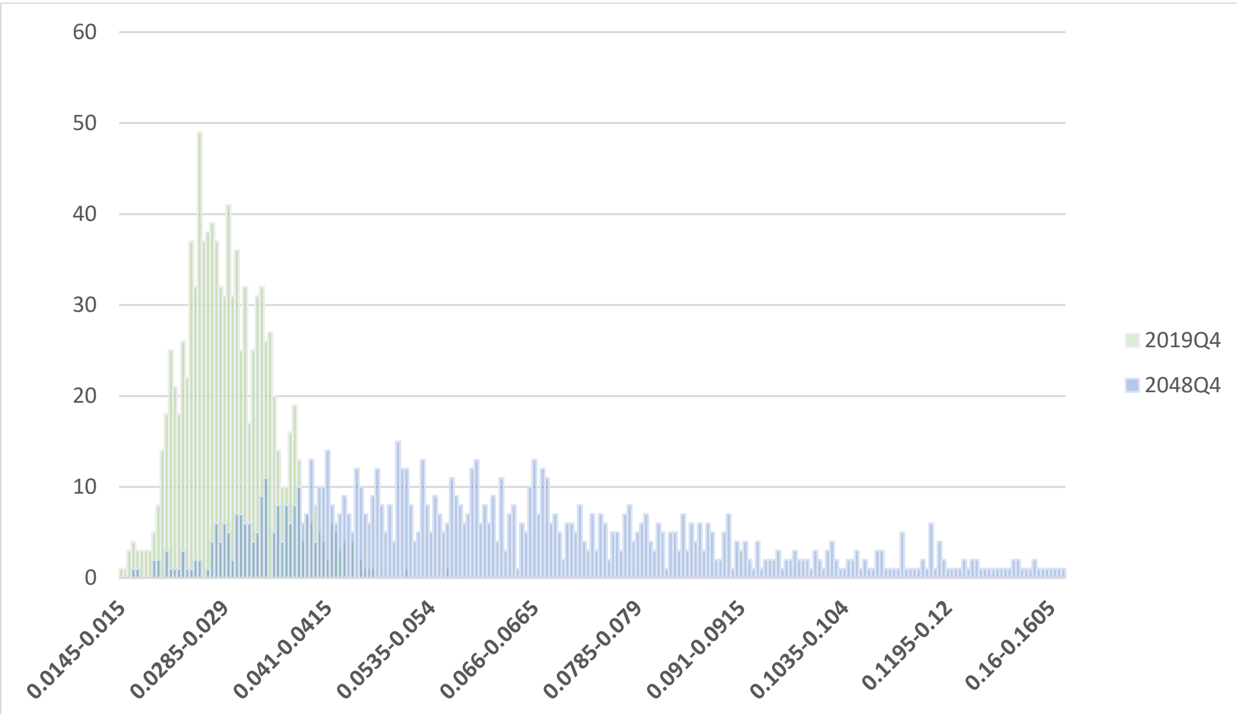

Target Summary

	2019 Q4	2048 Q4
Mean	0.0303	0.0627
Std dev	0.0055	0.0266

Simulation Summary

	2019 Q4	2048 Q4
Mean	0.0294	0.0636
Std dev	0.0055	0.0266

Prepared by Conning, Inc. Sources: ©2020 Bloomberg, L.P. and GEMS® Economic Scenario Generator scenario.

From this chart one can immediately compare the numerical values of the calibration targets and the simulated values of the calibration. Insight into the evolution of the distribution of the variable is also provided by producing histograms at different simulation horizons.

4.6.1 Target Setting

Calibrating and validating an ESG model involve setting targets for modeled variables that reflect an interpretation of history and a view on future trends of the variables. The targets summarize the expected behavior of simulated scenarios. Historical data are a natural and the most important input to the target-setting process. A long time series of historical data can show diverse evolution patterns of the modeled variables and shed light on future developments. The historical data used to establish the stylized facts to guide the model choice are usually combined with expert judgment to form the basis of setting targets.

Expert judgment is needed because a mechanical mechanism of target setting using only historical data can be problematic for various reasons. An immediate issue is the selection of an appropriate data window. A target setter needs to exercise judgment to choose a data window and to decide how much weight, if any, the model will place on more recent data relative to older data. For example, historical global inflation data show a universal trend of low inflation since the 1990s, a phenomenon understood to be the result of monetary policy regime change to inflation targeting by many central banks. Understanding inflation targeting and its implications for future inflation development will be helpful in setting inflation targets beyond what a deep statistical study of long history can provide.

As well as long-term targets, some users may need to set short-horizon targets (e.g., one year ahead). Setting shorter- or intermediate-term targets also needs careful judgment, especially when the targets are related to policy variables. For example, interpretation of changes in monetary policy stance may be the main determining factor for the short-term targets for interest rates. Uneven availability of historical data—referred to as ragged data—is commonly observed across different economies and market sectors. Blindly using historical data can lead to a structure of unrealistic targets. For example, some eastern European countries may not have the 1980s' high interest rates in their historical data but instead have extended periods of near-zero rates. If this leads to setting interest rate targets for those economies lower than those of neighboring economies with longer historical data, the targets may be unrealistic and misleading. Therefore, careful inter-economy comparison is indispensable to setting targets for a multi-economy ESG to ensure that a degree of consistency is achieved and to minimize any bias caused by data quantity or quality.

There is no single correct methodology for target setting, depending as it does on many factors. However, an idealized process might meet some or all of the following criteria:

- enables the setting of targets for key variables at a given horizon or horizons of interest;
- leads to targets that remain stable through time and do not require frequent revisions;
- is consistent across various economies;
- is not impeded by the lack of data availability;
- is justifiable based on the available data and validates well against relevant data; and
- reflects expectations of key stakeholders of the ESG.

Although satisfying all of these requirements may not be possible, they can serve as useful guidance in removing potential sources of observer bias.

4.6.2 Back-Testing

Back-testing can provide useful insight into the robustness of models, as one may have greater confidence in the robustness of the ESG if it can produce some scenarios that mimic extreme historical events.

Figure 4.14. U.S. year-over-year CPI inflation

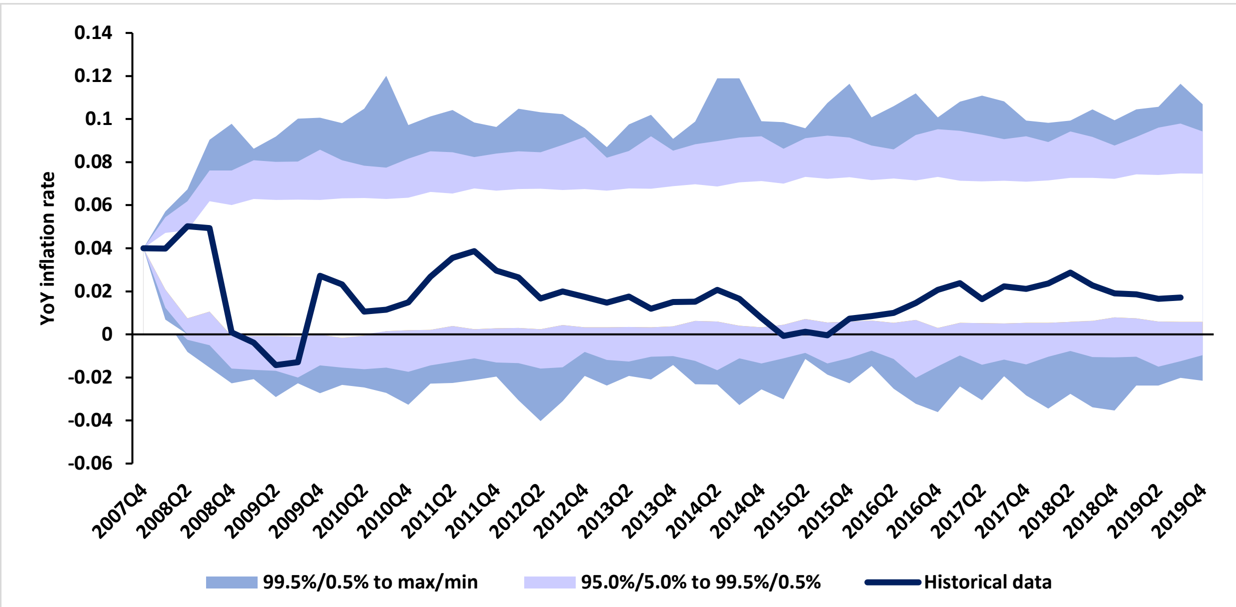

Prepared by Conning, Inc. Sources: ©2020 Bloomberg, L.P. and GEMS® Economic Scenario Generator scenario.
Note: YoY = year-over-year; CPI = consumer price index.

One can back-test any model if the model parameters for a certain period in the past are available. A successful back-test will result in the simulated model behavior incorporating what actually transpired in history out of sample. There is no specific statistical test that is normally done in this case; it is more a matter of establishing a level of comfort with what the model can produce.

4.7 Ongoing Maintenance

The process of ongoing maintenance of an ESG is based on the way in which the models, calibrations, and validation processes interact over time and what actions are taken given the model performance. In this respect, we have already touched on the following:

- The calibration brings the model performance consistent with market conditions.
- The choice of calibration targets reflects a view.

- The robustness of the scenarios depends on achieving calibration objectives, obtaining a range of scenarios that incorporate historical experience, and generating some extreme but plausible scenarios.
- The natural steps in preparing ESG scenarios are to calibrate the models and run a validation check.
- If validation raises issues, then a re-parameterization may be required.
- It may also happen that the limitations of the model dynamics are such that there is no parameterization that will achieve all validation objectives.

In a nutshell, the maintenance of an ESG operates in accordance with the diagram in Figure 4.15.

Figure 4.15. ESG maintenance

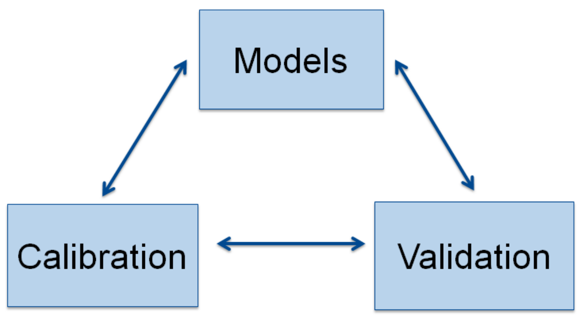

The most common and straightforward maintenance is when one updates initial conditions (i.e., market data) for each new calibration start date, but one does not re-parameterize all the parameters of the models within an ESG. When the ESG has been properly initialized to the current market data, a simulation run is made and a validation is performed. If some variables do not pass the validation criteria, a recalibration will be needed. When the recalibration is completed, another simulation and validation run is performed. This feedback loop is repeated until such time as the model passes all validation criteria.

A more severe and systematic failure of a model to perform adequately may require a more involved remedy. This can happen because changes in market conditions have left the original model dynamics significantly misspecified. In such cases, a new model may be required, and that will necessitate research carried out by specialists and may involve a considerable commitment of time and resources. For this reason, replacing an existing model is usually reserved for extreme cases and only for the most significant variables.

4.8 Summary

An ESG is a collection of models under a defined architectural structure that incorporates a specified level of detail and a selection of appropriate stochastic dynamics. Development and maintenance of an ESG calls for a careful approach to parameterization and a disciplined maintenance process to adequately reflect both historical and prospective financial dynamics.

ESGs are developed under a specific and intentional architecture that accommodates the appropriate interaction of component models. The way in which the component models are structured and interact with one another affects the causality and correlation structure of the ESG as well as the calibration methodology. This often employs a cascade structure and vector autoregressions. Interest rates and total returns are generally key output components.

When ESG simulation output is generated, an assumption needs to be made as to the level of detail to be included, including simulation frequency. ESG variables such as interest rates need to be computed and stored with a tenor structure. Price/income relationships, cash flow structures, prepayment features, and default events may also need to be incorporated. As a general rule of thumb, the greater the level of detail in an ESG, the slower it will run, and the larger the simulation data set that is stored.

Development is guided by the stylized facts and institutional details of the key economic variables to be modeled, as well as the level of detail needed for the application. Examples of the stochastic dynamics of a three-month T-bill and S&P daily return characteristics are illustrated, including pathwise characteristics such as jumps and volatility.

Data sources for economic modeling can vary in cost and availability. Examples of sources include Bloomberg, Thomson Reuters, Global Financial, Barclays Capital Live, central banks, bond-rating services, etc.

Parameterization/calibration is a process of selecting model parameters based on certain criteria. Estimating model parameters with historical data and calibrating models to specific market conditions are key parts of the process. The parameterization/calibration process and methodology often reflect a view, both explicit and implicit, with choices dependent on the application—business, regulatory, stress testing, investment management, etc.

The validation process involves checking that the calibrated model performs in line with the calibration criteria and that the general behavior of the model is consistent with the stylized facts. Checking that the model is performing in line with the calibration criteria will usually include a comparison of the simulated model statistics with the calibration targets. Back-testing can provide useful insight into the robustness of models.

The process of ongoing maintenance of an ESG is based on the way in which the models, calibrations, and validation processes interact over time. Under typical conditions, one updates initial conditions (i.e., market data) for each new simulation period, but one does not re-parameterize all the models of an ESG every period.

References

Bauer, M. D., and G. D. Rudebusch, "Monetary Policy Expectations at the Zero Lower Bound," Federal Reserve Bank of San Francisco Working Paper 2013-18, 2015.

Black, F., and M. Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81:3, 1973, pp. 637–654.

Chen, R.–R., and L. Scott, "Multi-Factor Cox–Ingersoll–Ross Models of the Term Structure: Estimates and Tests from a Kalman Filter Model," *Journal of Real Estate Finance and Economics* 27:2, 2003, pp. 143–172.

Christensen, J. H. E., "A Regime-Switching Model of the Yield Curve at the Zero Bound," Federal Reserve Bank of San Francisco Working Paper 2013-34, 2015.

Christensen, J. H. E, and G. D. Rudebusch, "Estimating Shadow-Rate Term Structure Models with Near-Zero Yields," *Journal of Financial Econometrics* 13:2, 2015, pp. 226–259.

Gouriéroux, C., and A. Monfort, *Simulation-Based Econometric Methods*, Oxford: Oxford University Press, 1997.

Hamilton, J. D., Time *Series Analysis*, Princeton, NJ: Princeton University Press, 1994.

Johannes, M., and N. Polson, "MCMC Methods for Continuous-Time Financial Econometrics," in *Handbook of Financial Econometrics: Applications,* vol. 2 in Handbooks in Finance series, edited by Y. Aït-Sahalia and L. P. Hansen, pp. 1–72, Amsterdam: Elsevier, 2010.

Robert, C. P., and G. Casella, *Monte Carlo Statistical Methods*, 2nd ed., New York: Springer, 2004.

Chapter 5: What Makes a Good ESG?

As we have previously discussed, ESGs can be used to generate real-world and/or risk-neutral economic scenarios. In this chapter we focus on real-world scenarios because that is the case of importance to casualty actuarial practice. Real-world scenarios are concerned with the future paths of economic variables, their potential influences and relationships, and the actual probabilities of various outcomes. It is this context that we must consider to better understand what makes a good ESG.

The question we address in this chapter is "What are the attributes of a good real-world economic scenario generator?" Given the many concepts connected with that question, there is no short answer to it. This chapter provides details on the concepts and associated criteria that are essential for a good ESG.

Let us think of an ESG as a software suite that produces a set of simulated economic scenarios where each simulated scenario represents a possible economic future. Taken together, the scenarios represent a distribution of future outcomes. The role of an ESG is to provide forward projections of economic and capital market variables in support of risk management and operational activities and strategies. In this context, a "good ESG" will be one that efficiently produces realistic economic scenarios with the level of detail required for the business use case.

Commonly, within an insurance organization, many ESG use cases will be operating simultaneously. For example, one team may be designing and pricing a new product, another team may be examining capital requirements, and a third team may be running a full enterprise risk management (ERM) program in which all aspects of risk are considered. Each of these teams may want to use different ESG calibrations to capture different aspects of risk. For example, the product design and pricing use case may require fewer ESG variables than the ERM use case, and the capital requirement use case may want to generate some extreme stress scenarios that are not as important for the other use cases. When one thinks about the role of an ESG in the broad context of an entire insurance organization, it is clear that a good ESG will have to simultaneously meet criteria across many use cases and be flexible enough to serve multiple roles. Here are some of the high-level features of a good ESG:

- It produces simulation results that reflect the economic view of the risk manager.
- Scenarios are consistent with realistic market dynamics.
- A large simulation should produce some extreme but plausible results (i.e., the simulation covers and moderately exceeds the benchmark stylized facts).
- Component models and architecture must have sufficient flexibility to serve in multiple roles.

If one discusses the essential features of a good ESG with a diverse group of ESG experts, those experts' lists of features and the relative importance of those features will vary. However, they will set forth a common core of ideas that can serve as a checklist of best practices. The following is a list and general discussion of those ideas.

A good ESG has a solid methodological foundation for the way the models are built and the way the variables are interrelated, and models are parsimonious, practical, and comprehensive. One characteristic of an ESG is that it is a collection of macroeconomic and financial models that fit or replicate key aspects of the historical data. A good ESG will construct those models with economic logic based on an understanding of the way key features of the financial markets relate to one another. For example, a good ESG will capture leverage effects in equity markets, and that requires that the dynamics of how volatility affects drawdown risk be built into the model. In contrast, an overparameterized econometric model that fits the data well may have poor simulation performance in part because it lacks stochastic dynamics that can produce the extreme but plausible scenarios that are critical in risk management. A good ESG strikes a balance between practicality and the ability to capture the critical features of the historical record. It is not practical to model every variable and every outcome precisely, and therefore one must determine which attributes are important to capture. One must also be pragmatic in which variables can be omitted, possibly leaving open the option for users to model variables of secondary importance on an as-needed basis. Models that are too simple will miss the true risks that affect an insurance business. Simple models do not capture the likelihood and scope of extreme events that could result in large capital charges. A good ESG does not "overfit" the data. Overfitting reduces a model's ability to explain or robustly simulate market behavior.

A good ESG provides a comprehensive suite of macroeconomic and financial variables and a multi-economy capability. For each individual economy, a good ESG will include models for the most important asset classes and economic variables that are characteristic of that economy. Some variables can be omitted subject to the guiding principle that variables of the greatest importance to the risk profile of a given firm are modeled. Most economies will include variables for

- sovereign interest rates and associated bond returns;
- equity index returns (often including dividend yield);
- inflation (possibly with specific sub-indices);
- GDP; and
- unemployment.

Many economies will include additional variables such as

- corporate bond yields and associated bond returns as well as transition and default characteristics (including possible differentiation of various corporate bond sectors);
- mortgage-backed securities;

- covered bonds;[5]
- municipal bonds; and
- interest rate, equity, and/or inflation derivatives.

Comprehensive coverage of asset classes and financial/economic variables is one of the vital considerations; the other vital feature is the ESG's ability to model globally correlated scenarios. Generating true multi-economy scenarios is much more than separately generating individual economies and tying them together after the fact; the correlation structures across economies should be an integrated part of the simulation process.

A multi-economy ESG simulates all of the single-economy variables and will also take into account the need to ensure realistic interactions between economies. For instance, there is significant co-movement in interest rates across many economies, and that needs to be captured. Other considerations include foreign exchange, correlations between the stock indices of different economies, and the dynamics of correlations that tend to increase during bear markets and weaken in "normal" times.

A good ESG also provides additional modeling capabilities such that the available asset classes and economic variables can be expanded in a consistent manner to meet more specialized user requirements. For example, a user may require a total-return index for an asset class that is not included in the ESG; additional modeling capabilities permit the user to construct such a total-return index using the existing ESG variables. The key feature is a mechanism for creating new asset classes that is integrated into the core ESG and can be built out as the role of the ESG evolves.

A good ESG can accommodate many types of calibration views across a wide range of benchmarks. An ESG needs to be able to reflect a wide range of economic views. One might calibrate an ESG to a view that embeds historical benchmarks for the past 50 years. Such a view would accept the statistical features of a 50-year data window as the calibration benchmark.

Another calibration view might consider short-term and long-term targets simultaneously, and those targets would recognize historical norms but would be set with some measure of expert judgment. Such a calibration view typically would consist of statements such as "One year into the simulation, the average 10-year Treasury yield should be 2.2%, and the long-run average 10-year Treasury yield should be 3.50%." Another user may require a calibration that reflects long-term historical dynamics but adjusts mean expectations to be consistent with current consensus opinions. Calibration views might embed externally mandated statistical targets that require minimum tail behavior on equity or fixed-income returns. This sort of calibration view is routinely required by the American Academy of Actuaries, the Canadian Institute of Actuaries, and regulatory authorities such as the Office of the Superintendent of Financial Institutions

[5] Covered bonds are debt instruments secured by a pool of mortgage loans or public-sector debt to which investors have a preferential claim in the event of default. For a detailed explanation of European covered bonds, see https://hypo.org/ecbc/covered-bonds/#essential-features-of-covered-bonds.

Canada. Yet another example of a calibration view would be to impose a range of interest rate behavior over the near term. For example, one might require that interest rates remain at low levels for the first two years of the simulation before gradually reverting to a higher longer-term level. In such a calibration view, one stops short of defining a known interest rate scenario but seeks to impose a strong requirement in order to test a specific business objective.

The key feature is that the ESG component models have effective and well-defined calibration methodologies, and preferably software has inbuilt tools for this purpose. The model must also be flexible enough to accommodate a wide range of calibration criteria for different uses. While accommodating this range of calibration criteria, the role of the ESG is to produce a range of simulations around those calibration views that captures the range of possibilities in future economic and financial paths based on a user's own views.

A good ESG produces simulation results that reflect a relevant view. When calibrated to reflect a relevant view, a good ESG will also produce simulation results that are relevant to historical facts. A common tendency is to overweight the experience of the recent past, but there is an important difference between setting a calibration view and then permitting realistic model dynamics to take over versus setting a calibration view that overwhelms any realistic market dynamics. The headline problems of the day have a strong tendency to dominate the modeling and thus the calibration process. This is understandable because users of an ESG are intensely focused on their business risks and those are generally aligned with the headline problems of the day. The danger in placing too much focus on headline risk is that one can forget that, over moderate time periods, the economy can migrate to new and different problems. For a risk manager and risk management application with a longer-term horizon, one must avoid the temptation to over-impose a view.

At the same time, a good ESG also facilitates stress testing, which, by design, is specifically set up to over-impose a view.

A good ESG produces some extreme but plausible outcomes. Global historical data teach us that the financial markets routinely experience various extreme events. The notion that a good ESG should produce some extreme but plausible scenarios is intuitive, but it is a broad concept rather than a blueprint. An *extreme* scenario is a scenario that reaches or surpasses the limits of past experience. Extreme scenarios are easy to imagine but many will not have a parallel in the historical record. A *plausible* scenario is a scenario that conforms to economic principles and that can be rationalized in an economic context. For example, a scenario for which long-term government bonds return more than equities or corporate bonds for an extended period of time is plausible, but a scenario for which inflation averages 20% and long-term Treasury yields average 10% is not plausible.

Recent interest rate behavior places a lot of strain on classical interest rate models. For example, experience with protracted periods of low interest rates has now established that the entire Treasury yield curve can remain low or negative for extended periods of time. This has been evident in Japan for many years and has also been a feature of the U.S. and European markets for the past several years. A good ESG should be able to produce scenarios that exhibit this feature while still allowing for scenarios with higher interest rates. However, it is difficult for these models to generate a 10-year period of low or negative interest rates that can then come back to life and resume a "normal" interest rate trajectory. New model features are probably needed to accomplish this.

One measure of the ability of an ESG to meet or exceed historical norms is to compare simulated returns with historical returns and observe whether the simulated returns have heavier tails than history. A good ESG generally will have this property, which one might think of as "covering history." Other concepts of extreme outcomes should be considered, depending on the use case.

A good ESG not only produces realistic distributions for key economic variables but also represents the correct "pathwise" behavior—how the simulation moves through time to produce a particular outcome. The pathwise behavior reflects the course of the variable in terms of periods of higher or lower volatility as well as jump behavior and volatility clustering. Pathwise behavior defines the way an insurance company or investor will experience the market evolution.

A good ESG embeds realistic market dynamics. Market dynamics are not merely statistical targets but attributes that reflect the way that a variable changes and relates to other financial variables through time. However, the idea that a good ESG should embed realistic market dynamics is a guiding principle rather than a specific modeling feature. Market dynamics vary greatly across time and economies, and no suite of ESG models can capture them all. A vast amount of research is done in this vein every year, and the conclusions often change over time. However, certain market dynamics stand out prominently in risk management and are essential for a good ESG. Important examples of market dynamics include the following:

- Yield curves can shift and twist in many directions, where a wide range of movements is possible.
- Over the longer term, there will be periods when one asset class outperforms followed by periods when that asset class lags.
- Over extended periods of time, the strength of the correlation of returns between asset classes changes.
- Corporate bonds may exhibit periods of fairly predictable migration patterns and then suddenly become unstable.
- Corporate bond spreads are subject to periods of extreme contraction and widening.

Following the standard practice, let us call a single simulated scenario of all modeled variables a "path," where a path represents one possible future evolution of the economy and therefore represents one possible complete future "economic experience." Taken together, a collection of simulated paths represents the simulation output from the ESG. When one looks at statistics such as average interest rates five years into the simulation, these are statistics that summarize cross-sectional behaviors across all simulated paths. Equally important are measures that relate to the attributes associated with individual simulated paths. Classic examples of such pathwise measures are fluctuations in correlation and volatility. A good ESG should produce simulated pathwise behavior that is faithful to market characteristics, because such a scenario is consistent with an evolution of the economy that an insurance company might experience and respond to with management decisions. For example, many return series have periods of both elevated and subdued volatility—a pathwise effect. If one considers only the distribution of returns, then one could be using a model that has an acceptable return distribution but fails to recognize that the volatility characteristics may not be reasonable. A good ESG should incorporate dynamics that produce pathwise fluctuations in realized volatility because they affect investment allocations, hedging costs and mark-to-market positions.

Realistic market dynamics involve a mixture of statistical targets, pathwise behavior, and a sufficient range of outcomes. Some of this performance can be measured using statistical validation tests, while other features are more qualitative.

A good ESG is computationally efficient and numerically stable. An ESG is a complex system of models, underlying variables, and parameters. The choice of models should take into consideration the properties that make them efficient to implement computationally, i.e., faster run times and greater accuracy. The choice of models should also be robust enough to maintain stable parameter estimates.

When calibrated to historical data windows, parameter estimates for the component models should exhibit some measure of robustness against changes in the data windows. The ESG component models should capture changing initial market conditions while exhibiting underlying parameters that evolve at a moderate pace. If model parameters that are calibrated to historical data fluctuate rapidly as the data window changes, that instability probably reflects significant model misspecification. Robust ESG models naturally tend to have stable parameter estimates as new data are added to a large historical data window.

Although there is an ongoing need to refine parameter estimates in light of evolving market risks, making significant parameter changes from period to period is disruptive to the broad ERM process because it would induce notable changes in the risk profile of various asset classes.

To be computationally efficient, a good ESG will take advantage of closed- or semi-closed-form pricing algorithms to compute prices and other variables whenever possible. ESGs that have no closed-form pricing algorithms can be cumbersome to estimate from market data, slow to simulate, and difficult to recalibrate to a user's views. A good ESG must make full use of efficient computational algorithms for pricing to ensure both good run-time performance and numerical accuracy.

A good ESG has fast and robust recalibration capabilities. In the post–financial crisis era, there remains considerable uncertainty about where the U.S. and other economies are headed and to what level key economic variables such as interest rates are likely to revert. Some practitioners even hold the view that it is unwise to assume interest rates will revert to any stable level; other practitioners try to synthesize the range of uncertainty into their economic scenarios by requiring considerable volatility in key variables.

Blending diverse possibilities into an ESG requires the capability to perform fast and robust recalibration to many different targets while continuing to initialize the models to their correct market values.

A good ESG meets the requirements of regulators and auditing firms. Given the heightened emphasis on risk and capital management within the insurance industry, a good ESG meets the growing demands of regulators and external parties such as auditing firms and rating agencies. ESG methodologies should be transparent to support regulatory requirements. This entails documentation of the maintenance processes (e.g., parameter and software updates) and other supporting materials that enable a comprehensive understanding of the theoretical and empirical integrity of the simulated scenarios.

Meeting the requirements of regulators and auditing firms does not mean that every idea put forward by those entities needs to be embraced, however. The ability to effectively reply to regulators and audit firms and to combine experience and expert judgment in addressing their concerns is important to maintain, either in house or on demand from a trusted consulting resource.

In more recent times, cybersecurity has also become an important part of software engineering. ESGs and the software platforms in which they reside must now be architected in a way that meets best practice standards and minimizes security risks.

A good ESG produces sufficient simulation detail for extensive validation. The process of validating an ESG is separate from model development but is part of the maintenance cycle that we discussed in Chapter 4. While it is self-evident that a good ESG validates to its calibration benchmarks, the process of performing the validation can require extensive inputs. A good ESG is capable of providing the very detailed simulation output that is needed to validate its performance. While one expects a good ESG to provide a range of total-return and yield outputs, one should also expect that a good ESG provides detailed outputs on default behavior, rating transitions, and the discount factors needed to perform regulatory risk-neutral testing.

5.1 Statistical Criteria

Statistical criteria are important in assessing the quality of an ESG. Statistical calibration criteria are usually numerically specified but can also be qualitative in nature. Statistical criteria belong to one of two broad categories: *qualitative features* and *quantitative measures*. The issues one must address in both categories are not amenable to a checklist approach, however, and expert judgment plays a role. This section will provide insights into each category and illustrate those with real market data.

5.1.1 Qualitative Features

Just as the stylized facts of financial markets guide the development of ESG models, so too do the stylized facts inform the quality of an ESG. Once the ESG component models are developed and a calibration is performed, an important first step in validation is to check that the most important qualitative stylized facts are satisfied by the simulated output. Examples of qualitative statistical criteria include these:

- the calibrated models produce sufficient mass in the tails of simulated distributions to capture or exceed historical extremes;
- the return-volatility characteristics of main asset classes tend to follow an orderly mean-variance frontier consistent with known market-risk characteristics;
- the simulated output produces some volatility clustering for variables that is typical of the market data;
- the simulated variables exhibit some stability of behavior within an appropriate range of possible future values;
- the interest rate variables have reasonable mean reversion properties—with plausible level and speed of reversion; and
- the co-movement and correlation between important economic variables are consistent with historical norms.

When checking how faithful an ESG is relative to statistical criteria, one usually compares the statistical properties of a calibrated simulation against a specific set of calibration criteria. A common specification of calibration criteria is to prescribe mean levels and volatilities at various simulation horizons. For example, the calibration criteria may require that the 10-year interest rate have an average value of 3% five years into the simulation. The percentiles of the variable at various simulation horizons might also be specified. While any aspect of the distribution could be specified, in practice it is typically the first few moments, the median, and extreme percentiles that tend to be of interest. These statistical quantities have tangible meaning, and the extreme percentiles are often used to define the regulatory or adequacy requirements of a simulation. A good ESG should have sufficient flexibility to be calibrated to statistical targets across several simulation horizon dates.

As a simple illustrative example, suppose that an ESG includes a one-factor interest rate model whose short-rate stochastic dynamics are described by a Cox–Ingersoll–Ross (CIR) process

$$dx_t = k(\theta - x_t)dt + \sigma\sqrt{x_t}dW_t.$$

This model admits a closed-form formula for the yield curve and the yield of maturity τ at horizon t is given by

$$y_t(\tau) = -\frac{1}{\tau}(A + Bx_t),$$

where

$$A = \frac{k\vartheta}{\sigma^2}\left\{(k+h)\tau + 2log\left[\frac{2h}{2h + (e^{-h\tau} - 1)(k+h)}\right]\right\},$$

$$B = \frac{2(e^{-h\tau} - 1)}{2h + (e^{-h\tau} - 1)(k+h)},$$

$$h = \sqrt{k^2 + 2\sigma^2}.$$

Suppose that this model has been estimated using historical data and calibrated to the current ($t = 0$) market conditions and that the estimates of parameters are given by

$$k = 0.2, \theta = 0.562, \sigma = 0.2,$$

$$x_0 = 0.005.$$

Assume that the theoretical yield curve based on these parameters fits perfectly the observed yield curve at time 0. Figure 5.1 depicts this yield curve.

Figure 5.1. Initial yield curve

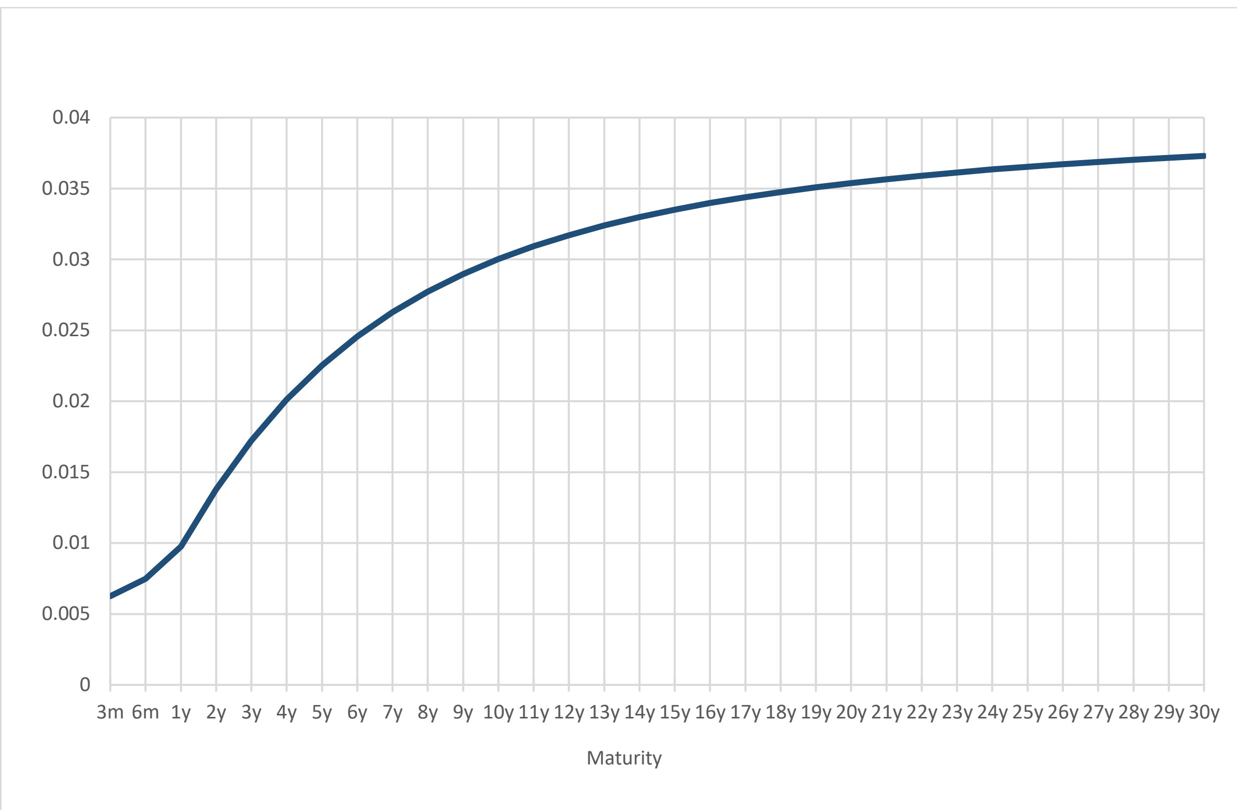

Prepared by Conning, Inc.

At a given future time horizon t, the mean yield for maturity τ can be calculated by substituting x_t with its theoretical expectation value

$$E[x_t] = x_0 e^{-kt} + \theta(1 - e^{-kt})$$

in the yield formula above. The chart in Figure 5.2 shows the path of the average 10-year yield based on the given parameter values. Now, suppose that targets for multiple simulation horizons are set as shown in the chart. The trajectory of the targets reflects the target setter's view that the 10-year yield will evolve to 5% rather than 4.8% on average in the long run and will reach that level around 15 years into the simulation.

Figure 5.2. Trajectories of 10-year yield averages and targets

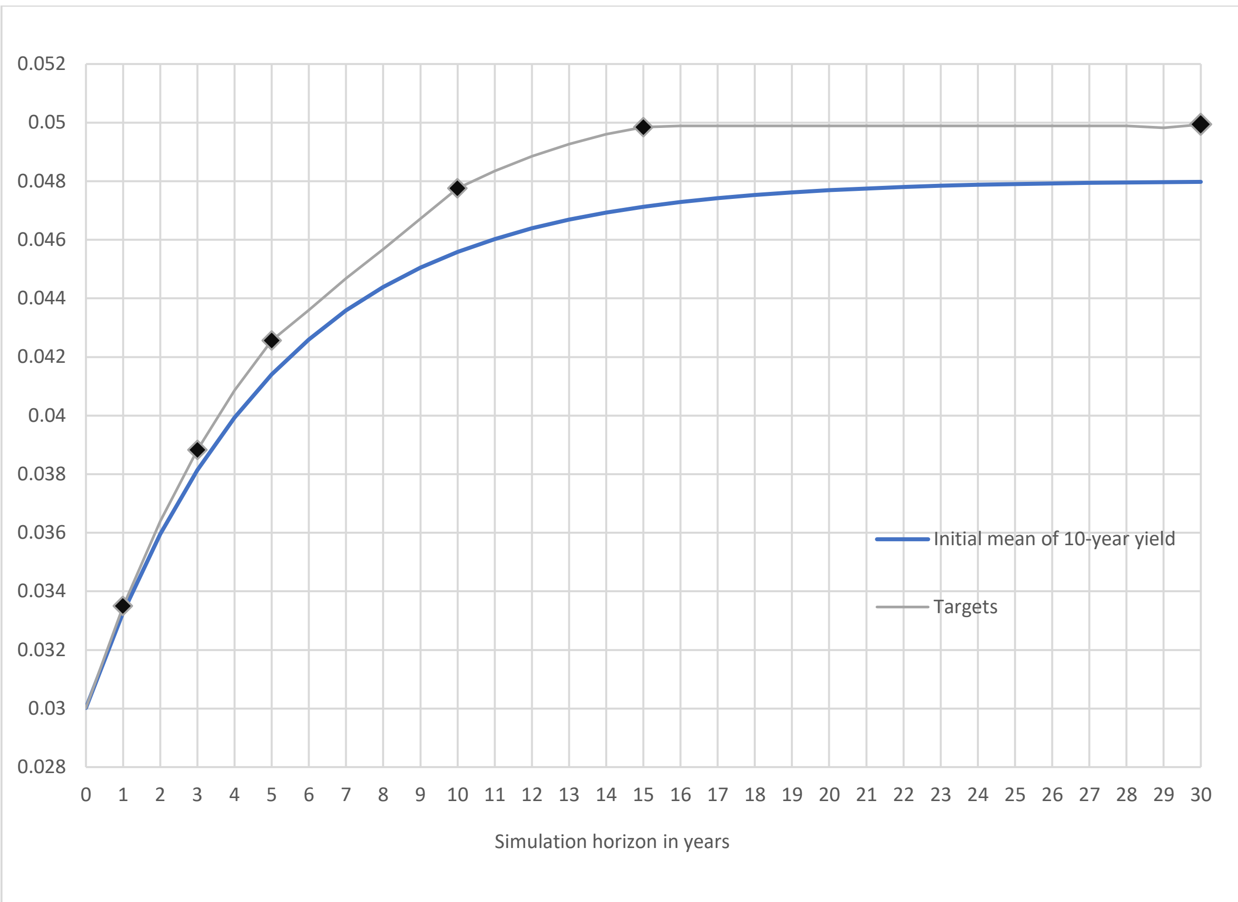

Prepared by Conning, Inc.

Calibrating the model to the targets can be described as an optimization problem in which we aim to find new parameter values that minimize the distance between the target trajectory and the trajectory of mean yields based on new parameter estimates while keeping the resulting theoretical initial yield curve close to the observed yield curve. The new parameter values obtained from this calibration are

$$k = 0.1987, \theta = 0.587, \sigma = 0.2,$$

$$x_0 = 0.0018.$$

The chart in Figure 5.3 compares trajectories of the initial average yield, new average yield, and targets, and it shows that most of the targets are hit closely with the new parameter values.

Figure 5.3. Trajectories of 10-year yield averages and targets

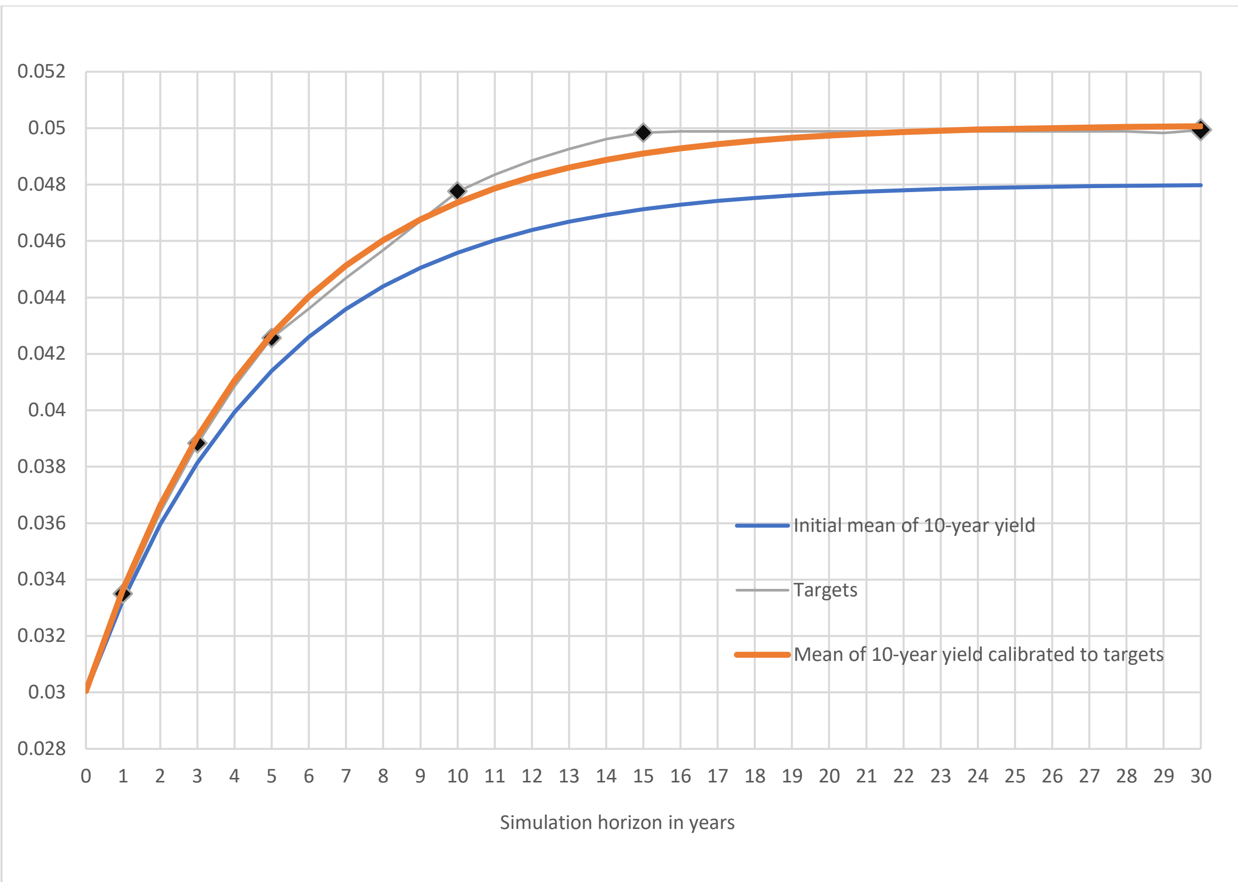

Prepared by Conning, Inc.

Figure 5.4 depicts the theoretical initial yield curve implied by the new parameters and the observed initial yield curve.

Figure 5.4. Initial yield curve

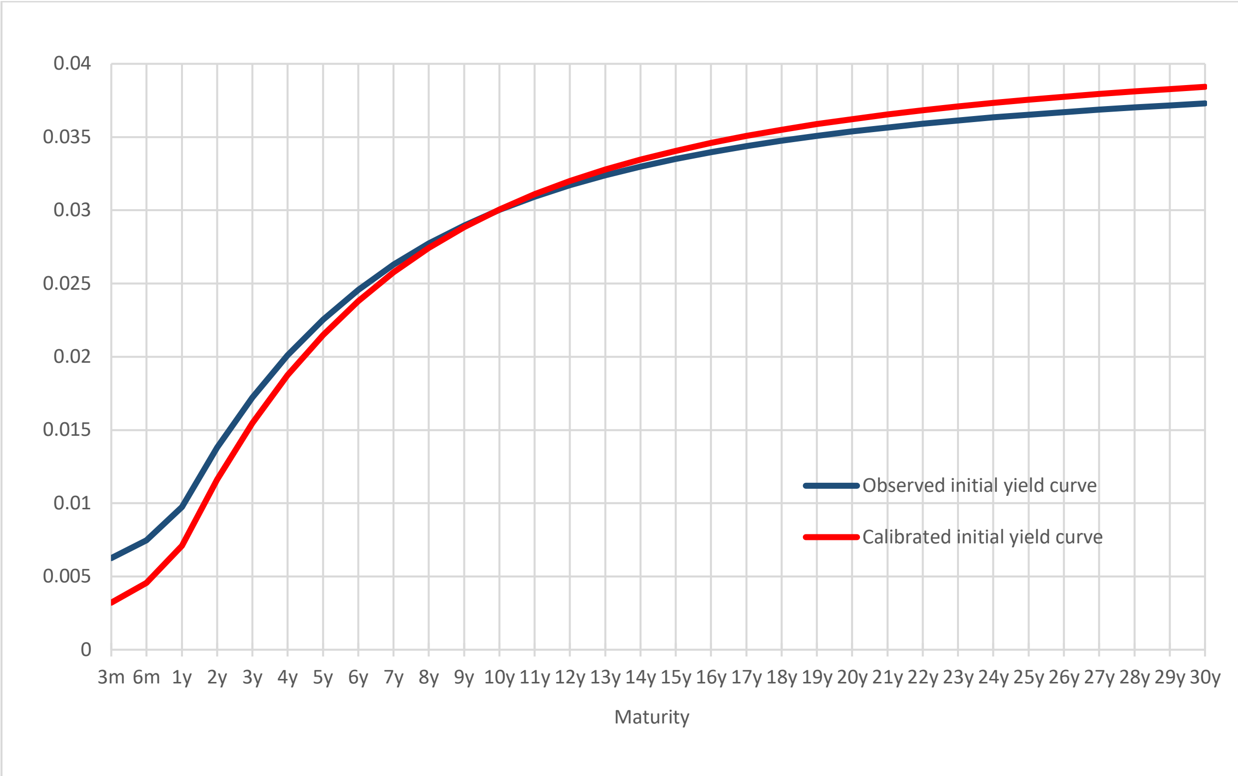

Prepared by Conning, Inc.

The above result hints at some of the limitations of the one-factor CIR model and in general of all the time-homogeneous short-rate models. When the model is calibrated to the targets, the initial yield curve from the simulation will deviate from the observed curve. An approach to handling this issue is to transform the endogenous model into an exogenous model by including deterministic time-varying parameters.[6] The one-factor model, even with a time-dependent extension to fit the initial term structure, is still not flexible enough to accommodate other kinds of targets. In the above discussion, targets are set for the yield of a single maturity—that is, the 10-year yield—but it would be more realistic to set targets for multiple maturities at several simulation horizons, which is practically equivalent to targeting the shapes of yield curves at multiple time horizons. More complicated and qualitative targets may include the tail property of bond returns implied by the yield curve evolution or skewness of the distributions of yields at various time horizons, etc. Adding more factors to the interest rate model will be necessary to address complex systems of targets like these.[7]

[6] See Brigo and Mercurio (2007) and the discussion later in this *User's Guide* in Section 6.3 for a more detailed explanation of the classical models and their extension.

[7] See Section 6.3 and literature therein for a discussion of the number of explanatory factors in an interest rate model.

In setting the statistical targets, it is important to note that the targets should be "coherent." In the following we give examples of statistical targets that are coherent and those that are not; the distinguishing feature between the two is that coherent statistical targets fit naturally with the way the economic variables behave. With experience, one develops an operational feel and understanding for the types of targets a good ESG can achieve.

Many important financial variables tend to have average levels that rise or fall over time in a smooth pattern. The average levels for such variables will have a smooth concave pattern when they are stationary, which is a common assumption for modeling various interest rates. The rate at which the average level changes over time is governed by the mean reversion intensity of a model. The chart shown in Figure 5.5 is typical of a stationary variable.

Figure 5.5. Targeting the average trajectory for an ESG variable

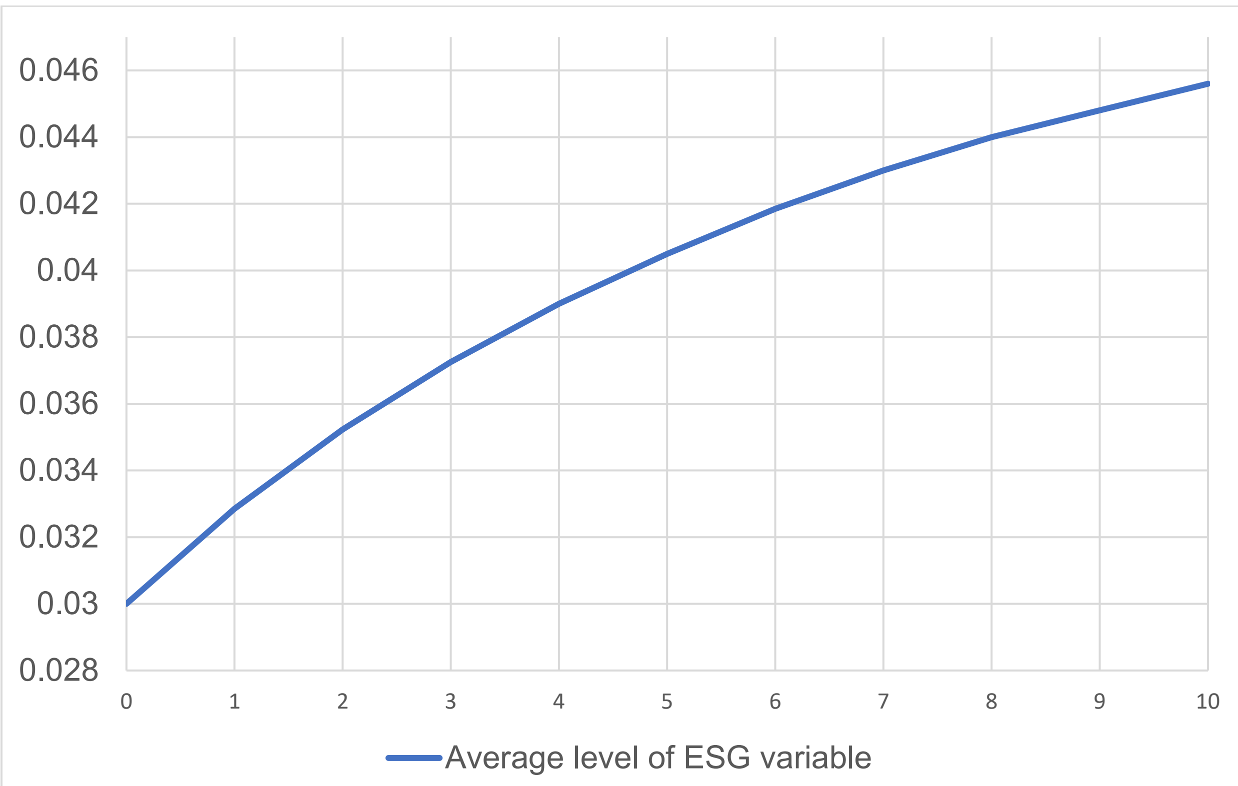

Prepared by Conning, Inc.

We might consider another example in which the targets we set lead to problematic model dynamics. In this case the variable starts at a value of 0.03 (3%) and on average rises toward a higher long-run level over the 10-year simulation horizon. The curvature of the average-level curve is characteristic of a mean-reverting process, for which the movement toward the long-run level is larger the farther the variable is from its long-run level. Now, suppose that we have a set of calibration targets for this variable, as shown by the red dots on the graph in Figure 5.6.

Figure 5.6. Targeting the average trajectory for an ESG Variable

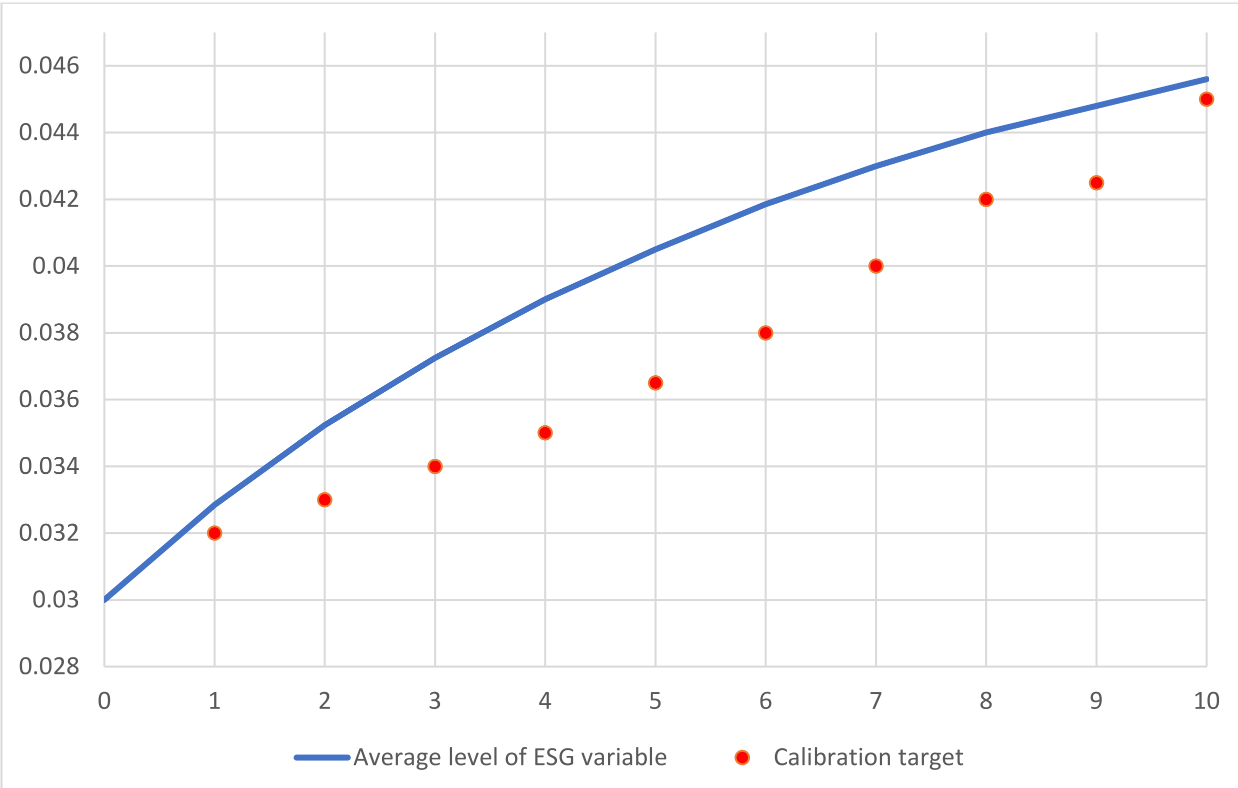

Prepared by Conning, Inc.

To calibrate to those targets, we cannot change the starting value of the ESG variable at 0.03, since initial values are typically set by market conditions. However, some model parameters can be changed to improve the calibration fit, and the basic approach is to adjust the mean reversion parameters and mean reversion levels. Let us suppose that Figure 5.7 shows the result of this adjustment.

Figure 5.7. Targeting the average trajectory for an ESG variable

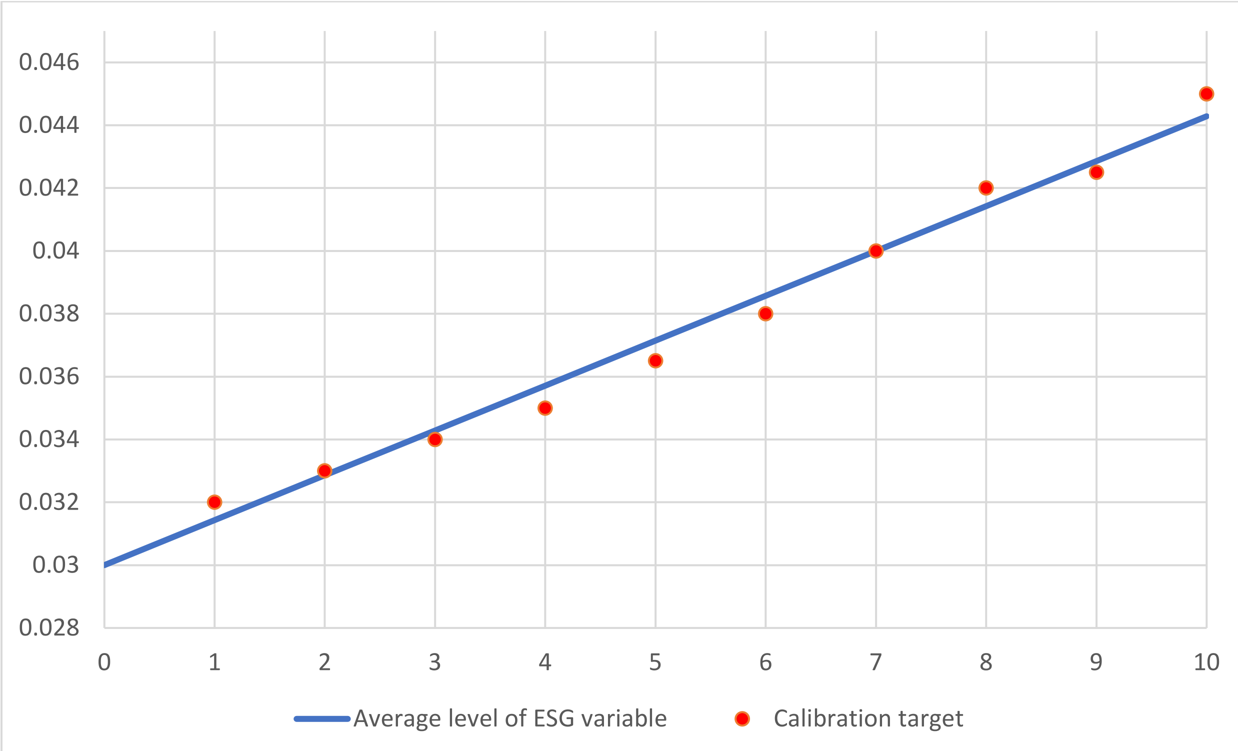

Prepared by Conning, Inc.

While the average levels are fairly close to the calibration targets, the calibration has some undesirable properties. First, the targets are linear in nature, which forces the model to have a very low mean reversion rate in order to produce a linear average level. With proper mean reversion, the average levels would have some curvature and could not come close to the targets. Also, the targets require the variable to rise 150 basis points on average over the 10 years, which is not unreasonable. However, the linear form of the average-level curve implies that this variable will keep rising on average for longer simulation horizons. The calibration targets given are not natural or consistent with the way in which interest rates and other economic variables tend to behave empirically. While this model can capture such calibration targets, the calibrated model performance may be poor, and one should not assess the quality of this model relative to such a set of calibration targets. Let us consider a revised set of calibration targets for the model as shown in Figure 5.8.

Figure 5.8. Targeting the average trajectory for an ESG variable

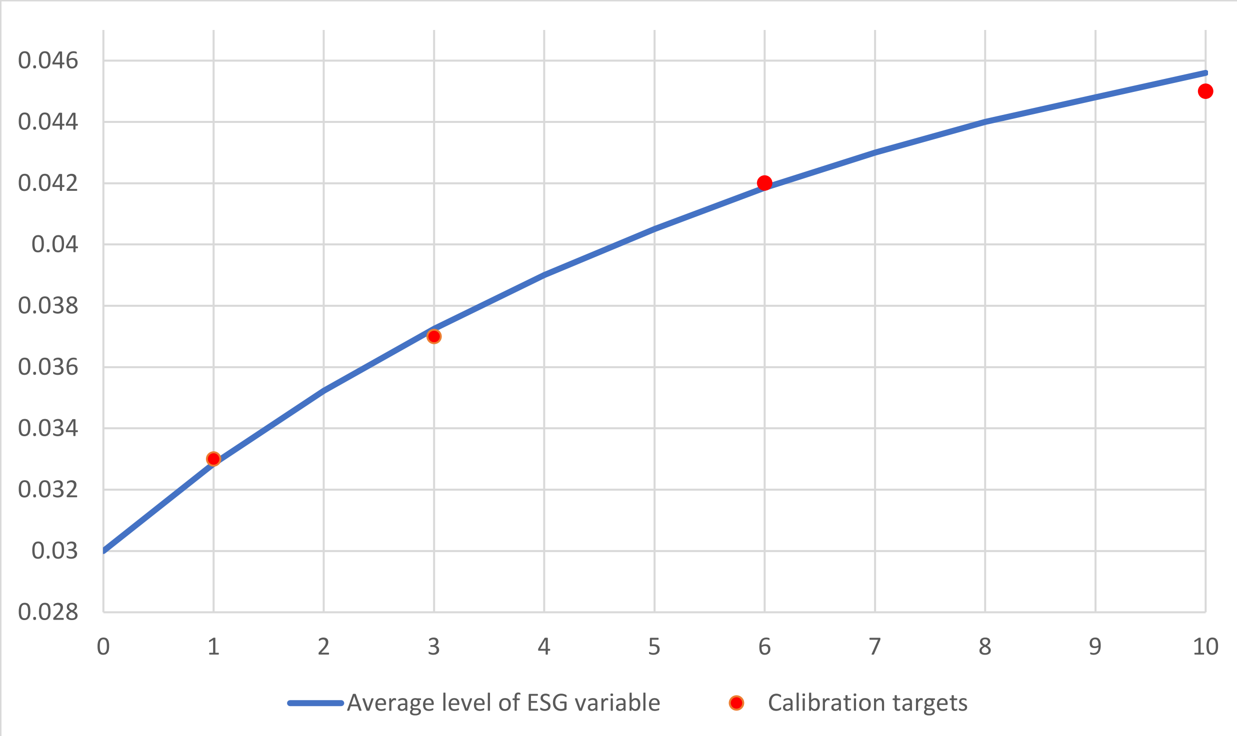

Prepared by Conning, Inc.

These targets have a fairly smooth behavior that is consistent with mean levels under a mean reversion dynamic. It is possible to achieve a good calibration for the model, and the long-run simulation performance of this calibration is likely to be stable. A good ESG component model should be able to calibrate to mean-level targets such as these.

The rate of change in the evolution of the mean simulated value of a variable is used to measure the rate of mean reversion in a calibration. If the mean simulated value slowly changes and moves toward a long-run value, then one considers that calibration to have a slow rate of mean reversion. A good ESG will provide some control over the rate of mean reversion for important variables such as interest rates.

There are many other qualitative statistical features that one might require of a good ESG. For example, one might want to capture the tendency of corporate bond spreads to compress as Treasury yields rise and to widen as Treasury yields fall, particularly if the ESG is being used to test dynamic investment strategies. The spread compression property is illustrated in the chart for U.S. A-rated 10-year bond spreads shown in Figure 5.9. The red line denotes the average response (measured as a regression line) of the change in A spreads to a change in Treasury yield levels.

Figure 5.9. Compression property of U.S. A spreads

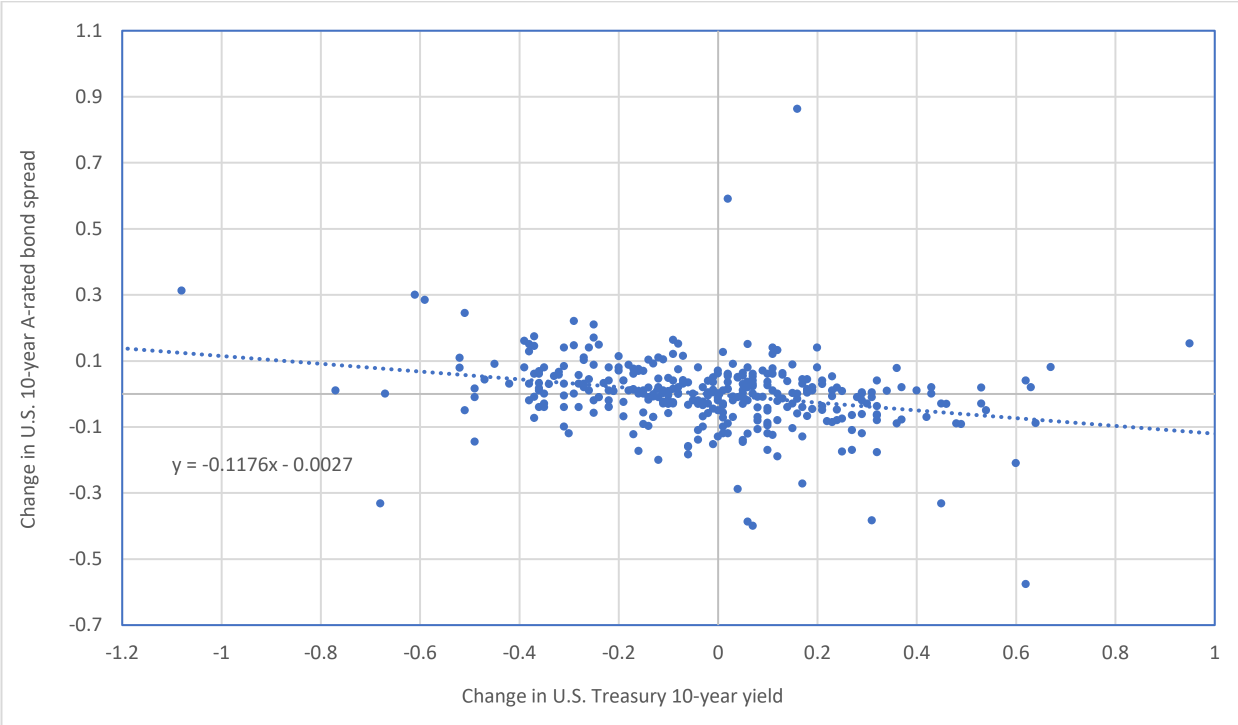

Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.

It would be difficult to precisely quantify what this feature should look like in an ESG. Capturing the qualitative behavior is the best one can do with a reasonable assumption about the response of spreads to changes in Treasury rates.

5.1.2 Quantitative Features

The ability to specify detailed quantitative targets and calibrate an ESG to them is fundamental. A good ESG will be capable of being calibrated to coherent targets across multiple simulation horizons. The specification of targets might be as simple as those shown in Table 5.1.

Table 5.1. Calibration targets

Simulation Horizon (Years)	Average Level
1	0.033
3	0.037
6	0.042
10	0.045

Tabular calibration targets will usually include targets for mean levels and volatilities. The target values can be guided by a view on the economy, benchmarks from historical data, or regulatory calibration requirements. A good ESG will have the flexibility to accommodate a broad range of such calibration criteria.

5.2 Pathwise Criteria

A path represents one possible future evolution of the economy and therefore represents one possible complete future "economic experience." A simulated set of economic scenarios for a given variable can be thought of as constituting a matrix of values, where the entries in a given row are the simulated values of that variable along a specific path, and the entries in a given column are the simulated values of that variable at a specific simulation time across all paths. Extracting a column of values from this matrix corresponds to the (cross-sectional) distribution of the variable at a given simulation horizon, and it is this type of distribution that is the focus of most statistical criteria. If we extract a row of data from this matrix, then we are looking at the evolution of the values of this variable for one simulation path across all simulation time. It does not make sense to specify calibration criteria as statistics for a specific path. For example, one would not require the average value of the 10-year interest rate for path 1 to be equal to 2%. In contrast, one might require the average value of the 10-year interest rate for all paths to be equal to 2%.

The importance of pathwise model behavior is that it is the simulated path that represents the way an insurance company will experience the evolution of the economy. In other words. the pathwise behavior is the only thing of interest when we want to investigate simulation dynamics. If the overall distribution of returns for an asset class is correct but the pathwise behavior does not correspond to the nature of the fluctuations that we see in the historical record, then there is a potential model issue.

Some financial variables can exhibit sudden changes in value that persist for a relatively short period of time before returning to more typical levels. Credit spreads provide an instructive example of the importance of pathwise behavior. The chart in Figure 5.10 shows U.S. 10-year investment grade corporate bond spreads.

Figure 5.10. U.S. 10-year investment grade bond spreads

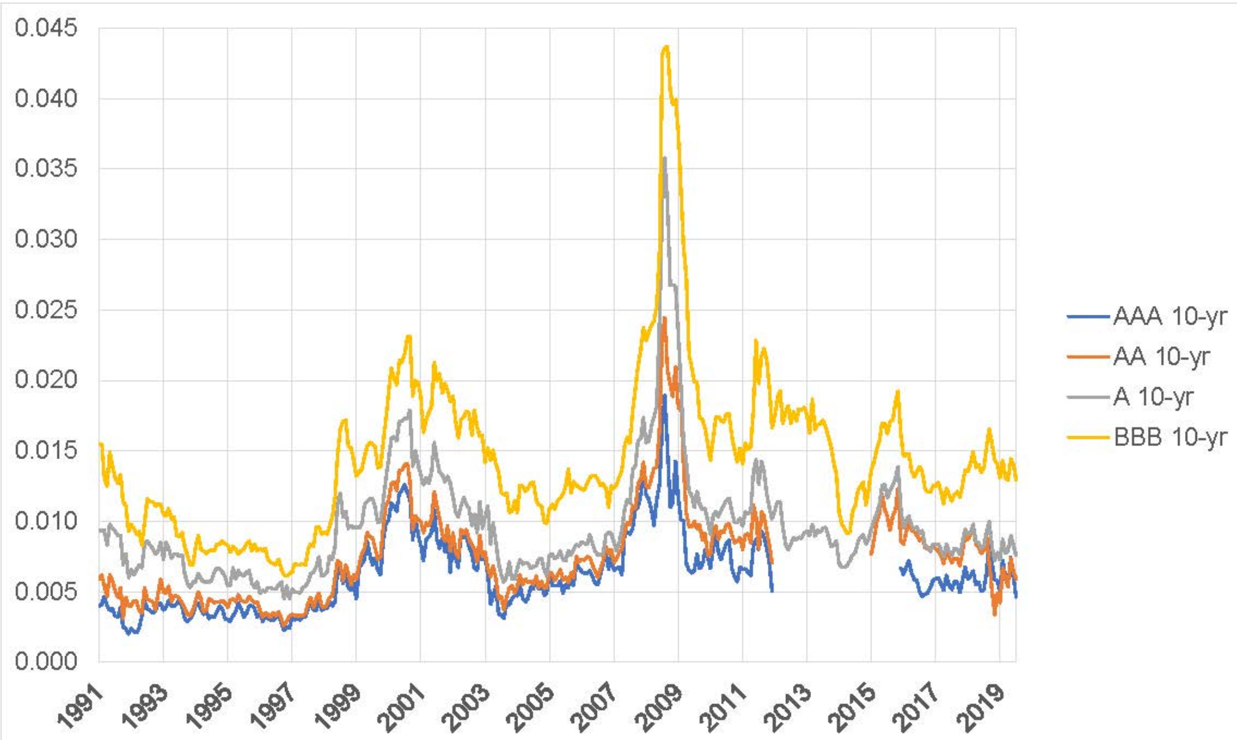

Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.

We see that higher-credit-quality bonds should normally have lower spreads, so that the credit spread curves should never cross each other. We also see that credit spreads can suddenly spike. Whether such spikes are jumps is less important than the fact that a good credit model should produce some simulated paths with spread "blowouts." Such spread blowouts will cause significant mark-to-market losses in corporate bonds, even if the event is short-lived and passes with limited defaults. The critical feature we want a model of credit spreads to have is the possibility of spread blowouts, meaning sudden explosions in credit spreads that abate over time. We probably do not want to be too specific about how long such blowouts last or precisely what level the blowout reaches; the important thing is that the model provides for this qualitative feature along some simulated paths.

As we saw in Chapter 4, equity volatility exhibits considerable clustering. Equity return distributions are also heavy tailed. A good ESG component model for equity returns will capture the fact that the distribution of returns is heavy tailed, but a good ESG equity model should also capture the pathwise features of volatility. A regime-switching equity model can do a good job of capturing heavy-tailed return distributions, but while the cross-sectional distributions can exhibit heavy-tailed behavior, each simulated path exhibits constant volatility when there is no change in regime and therefore fails to capture the volatility clustering that is characteristic of equity returns. The model fails to capture the way in which an investor actually experiences equity returns and should be considered as a deficient model for real-world simulations. The pathwise features of the model are determined by the stochastic dynamics selected in its design. Regime-switching equity models might be ruled out by a qualitative criterion on pathwise behavior.

The last example of pathwise behavior we will discuss regards the need to consider the co-movement of financial variables. While correlation is an important way to measure co-movement, we are also concerned with the broader idea that some ESG variables need to generally move together. A classic example is inflation and the short-term interest rate, as shown in Figure 5.11.

Figure 5.11. U.S. inflation and the short-term interest rate

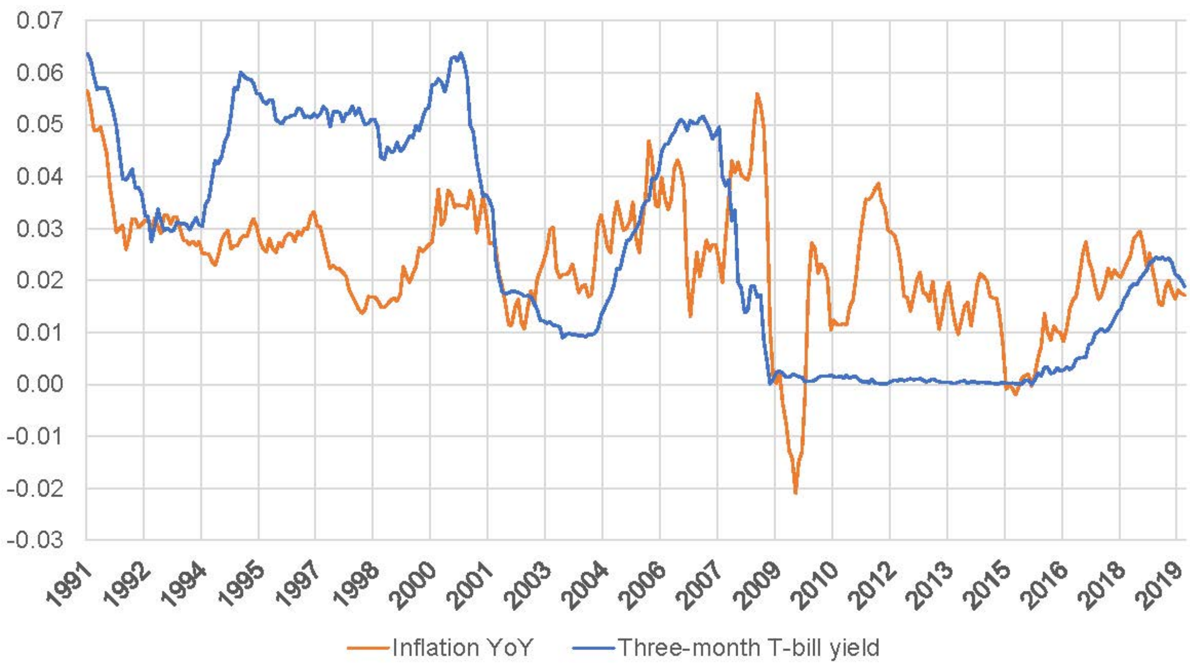

Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.
Note: YoY = year over year.

Inflation, measured in the chart as a year-over-year rate, and the short-term interest rate vary together but do not move in lockstep. The two variables should not drift too far apart from one another as they rise and fall. One might consider these variables to be cointegrated in the sense of the econometric term. Operationally, the important point is that we would not be satisfied with an ESG for which inflation and the short-term interest rate differed by a large amount for an extended period of time. A good ESG ensures that the co-movement of variables reflects historical market behavior.

5.3 Real-World Validation Considerations and Examples

The formal validation of ESG output is a crucial step in assessing the quality of an ESG. The fundamental process for real-world validation involves comparing calibration criteria against simulated model performance, and the criteria used are both qualitative and quantitative. Since expert judgment figures considerably in establishing calibration benchmarks, it is not surprising that considerable expert judgment is needed in validating an ESG. Qualitative features play a vital role in addition to quantitative measurements.

Concerning qualitative criteria, one is generally validating the ESG output against the qualitative features we have previously discussed for statistical and pathwise characteristics. In practice, one must decide which of those characteristics are the most important. Often not all of them can be accommodated by even very sophisticated ESGs. The process of understanding which qualitative model characteristics a good ESG needs and which are most important to the risk management process will then lead to the qualitative validation criteria that are to be applied. Validation is then the process of assessing how such criteria compare with the simulated output.

5.3.1 Quantitative Validation Checks

The most basic validation checks are purely quantitative and involve checking that the calibrated simulation has hit its calibration targets. A simple method for this validation step is to use a tabular summary. If one is checking that target levels have been hit at various simulation horizons, one might have a simple report such as the one shown in Table 5.2

Table 5.2. Simple tabular summary

Target Summary

	2019 Q4	2048 Q4
Mean	0.0303	0.0627
Std dev	0.0055	0.0266

Simulation Summary

	2019 Q4	2048 Q4
Mean	0.0294	0.0636
Std dev	0.0055	0.0266

Generally, one has a tolerance in mind to use as a guiding principle as to whether the simulation performance is sufficiently close to the calibration targets. Fan charts such as that shown in Figure 5.12 can also be useful in evaluating the success of a calibration.

Figure 5.12. Sample fan chart

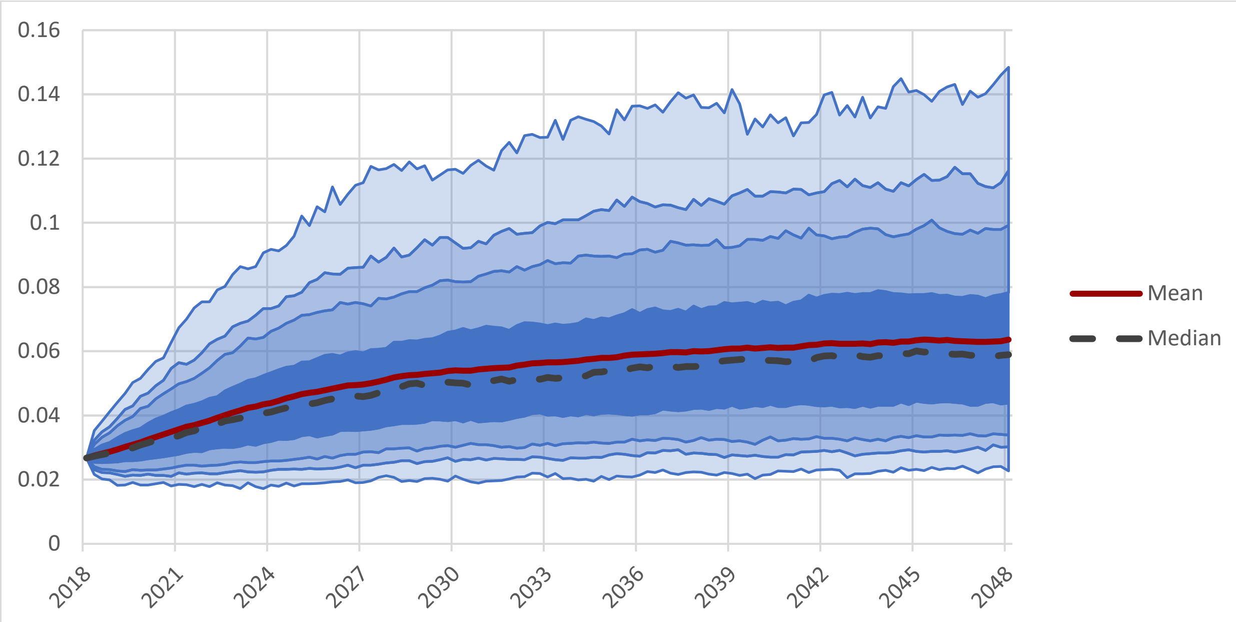

Percentiles plotted: 1%, 5%, 10%, 25%, 50%, 75%, 90%, 95%, and 99%. Prepared by Conning, Inc. Source: GEMS® Economic Scenario Generator scenario.

More detailed tabular validation checks are often performed, where all calibration targets are compared against their appropriate simulation statistics. Sometimes a tabular report is generated for review. The process of checking calibration targets against calibrated simulation results is easily automated so long as one can formulate an acceptance/rejection criterion that one is comfortable with.

Qualitative real-world validation checks are more subtle and are best discussed in the context of specific examples.

5.3.2 Checking Whether a Calibration Covers Historical Extremes

The process of checking that the tails of simulated distributions have sufficient mass to capture historical extremes is usually done using a histogram. In Figure 5.13, it would be reasonable to conclude that the calibrated simulation is adequate for capturing historical extremes.

Figure 5.13. Histogram

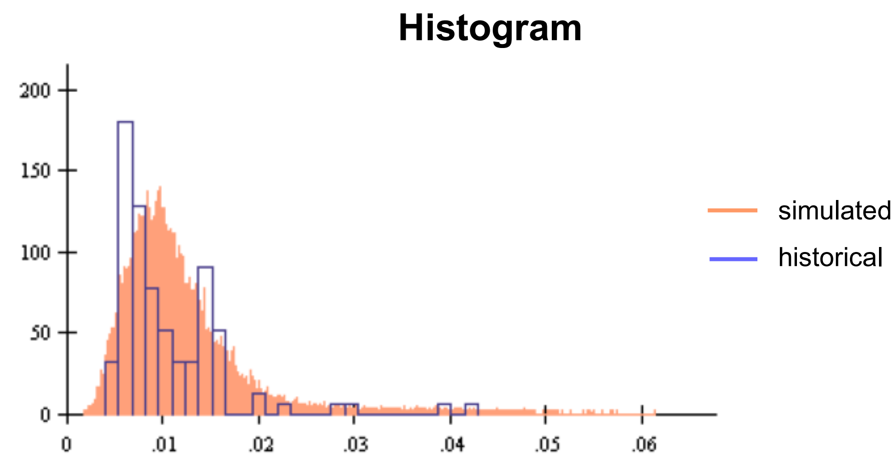

Prepared by Conning, Inc.

The nature of this check is visual. Of course, one can also quantify the simulated tail mass and perform numerical comparisons of simulated tail mass versus empirical estimates of historical tail mass. However, since a lot of expert judgment goes into quantifying estimates of historical tail mass, this would not be a precise exercise. For this reason, many users will be satisfied with a general qualitative check.

5.3.3 Checking Risk–Return Consistency across Asset Classes

A fundamental principle of finance is that the riskier an asset class is, the higher should be its expected return. A useful check on a calibration is to examine a *mean-variance*, or *risk–return*, *plot* of the important asset classes to see if they plot in a reasonable fashion. An irregular mean-variance plot can indicate that relative returns may need adjustments. Getting a regular mean-variance plot for major asset classes is important when using an ESG calibration for strategic asset allocation and other investment management processes.

Figure 5.14 shows an illustrative mean-variance plot for some simulated asset classes.

Figure 5.14. Risk–return profile for illustrative asset classes

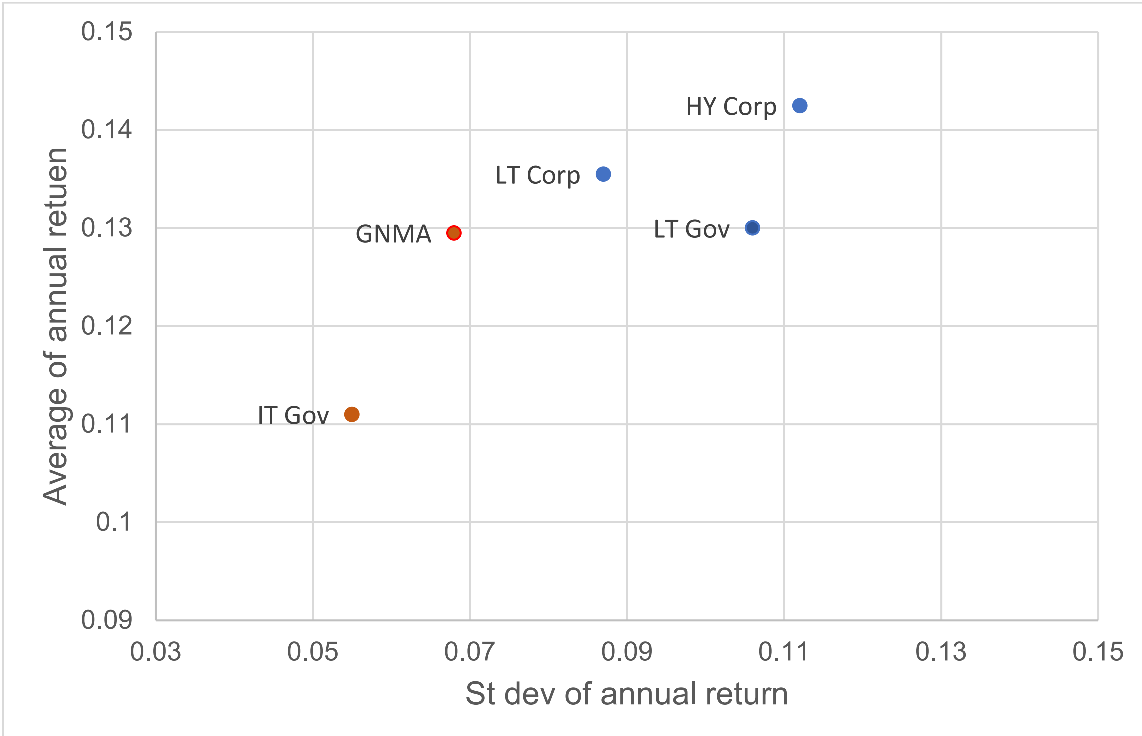

Prepared by Conning, Inc.

As a general principle, when an asset class plots to the "northwest" of another asset class, then the mean variance of that asset class dominates its neighbor. In an investment optimization setting, mean-variance-dominated assets are usually not selected for investment. In this chart, the asset class labeled LT Gov stands out as having undesirable risk–return characteristics relative to the other investment opportunities, and this might indicate a problem with the calibration. It is possible there are other problems as well, but to assess this we need to know more about each of the asset classes being plotted. While such irregular mean-variance plots are common when plotting asset classes from historical data, regular relative risk–return characteristics are important in many ESG use cases.

To make sense of such a plot, one needs to have an understanding of the relative risk–return characteristics of the asset classes being studied. In this example, there are five asset classes:

- Intermediate-term government bond returns (IT Gov)
- Long-term government bond returns (LT Gov)
- Long-term corporate bond returns (LT Corp)
- High-yield corporate bond returns (HY Corp)
- Government National Mortgage Association mortgage-backed security returns (GNMA)

Among the five asset classes, one might expect IT Gov to have the lowest risk and lowest return and HY Corp to have the greatest risk and greatest return. LT Corp might be a bit riskier than LT Gov and therefore have a correspondingly marginally higher return. GNMA might be considered less risky than LT Gov with similar to slightly lower returns. Therefore, we might expect the calibrated mean-variance plot to look more like the chart in Figure 5.15.

Figure 5.15. Risk–return profile for illustrative asset classes

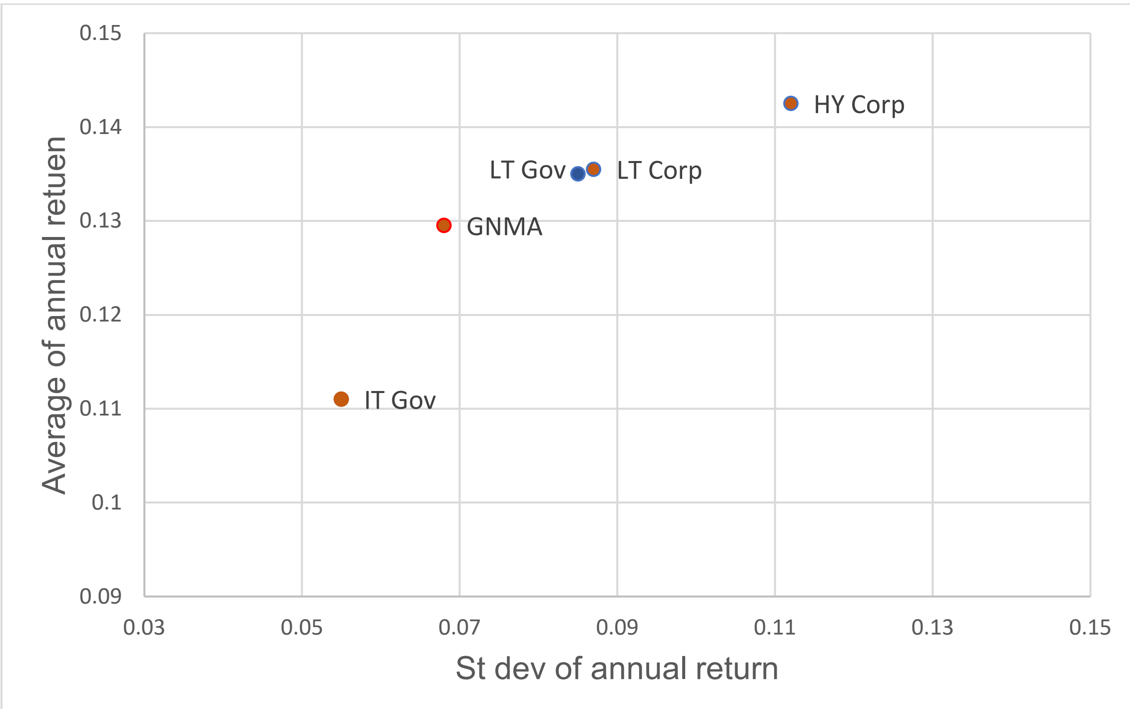

Prepared by Conning, Inc.

To arrive at the relative risk–return relationship in this chart, one would have to recalibrate some of the ESG component models. The validation process would flag the original calibration as having undesirable risk–return characteristics and suggest a corrective recalibration.

5.3.4 Check on MBS Model—Negative Convexity and Relationship to Treasuries

Mortgage-backed securities (MBSs) are an important investment class for property/casualty insurers, and negative convexity is an important feature of MBS price behavior. Negative convexity will show up when plotting MBS prices versus refinancing incentive. In Figure 5.16, MBS prices are shown on the vertical axis, and the refinancing incentive (difference between mortgage rate and current refinancing rate) is shown on the horizontal axis.

Figure 5.16. Market price and fitted model prices as spread varies

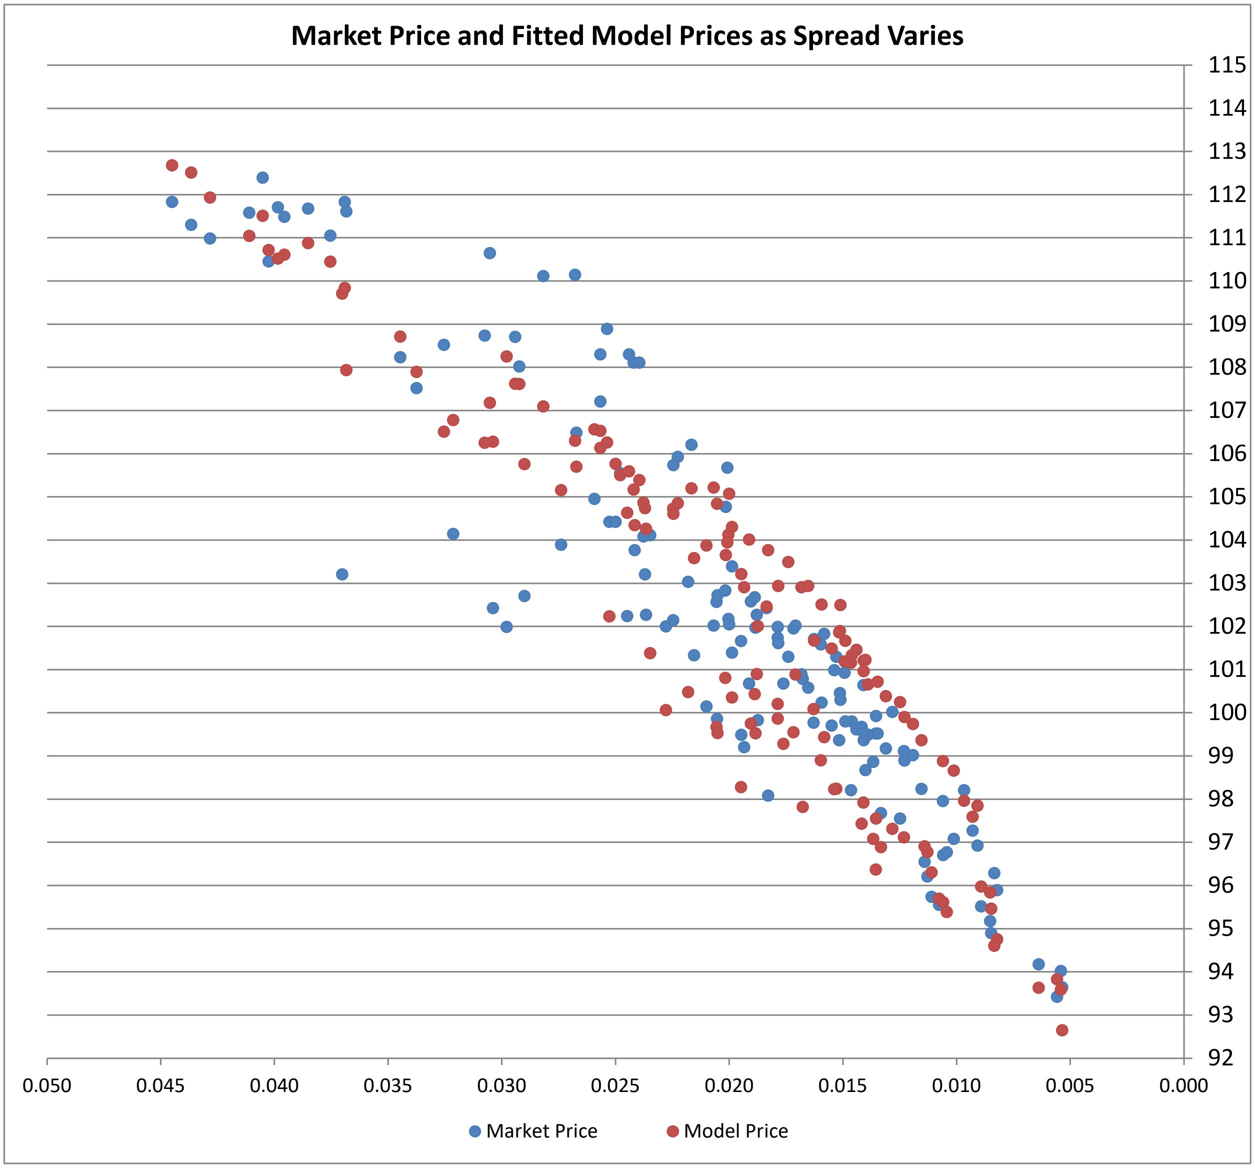

Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.

The negative convexity is evident in the market data and the ESG component model prices. Some ESGs do not explicitly model MBS prices and instead model MBS total returns directly, perhaps because modeling MBS prices is quite intricate. If an MBS total-return model is correctly calibrated, one should see aspects of negative convexity in a plot of MBS returns versus Treasury returns, as we see in the chart based on historical data in Figure 5.17.

Figure 5.17. Monthly GNMA total return versus U.S. Treasury total return (Jan. 1976–June 2019)

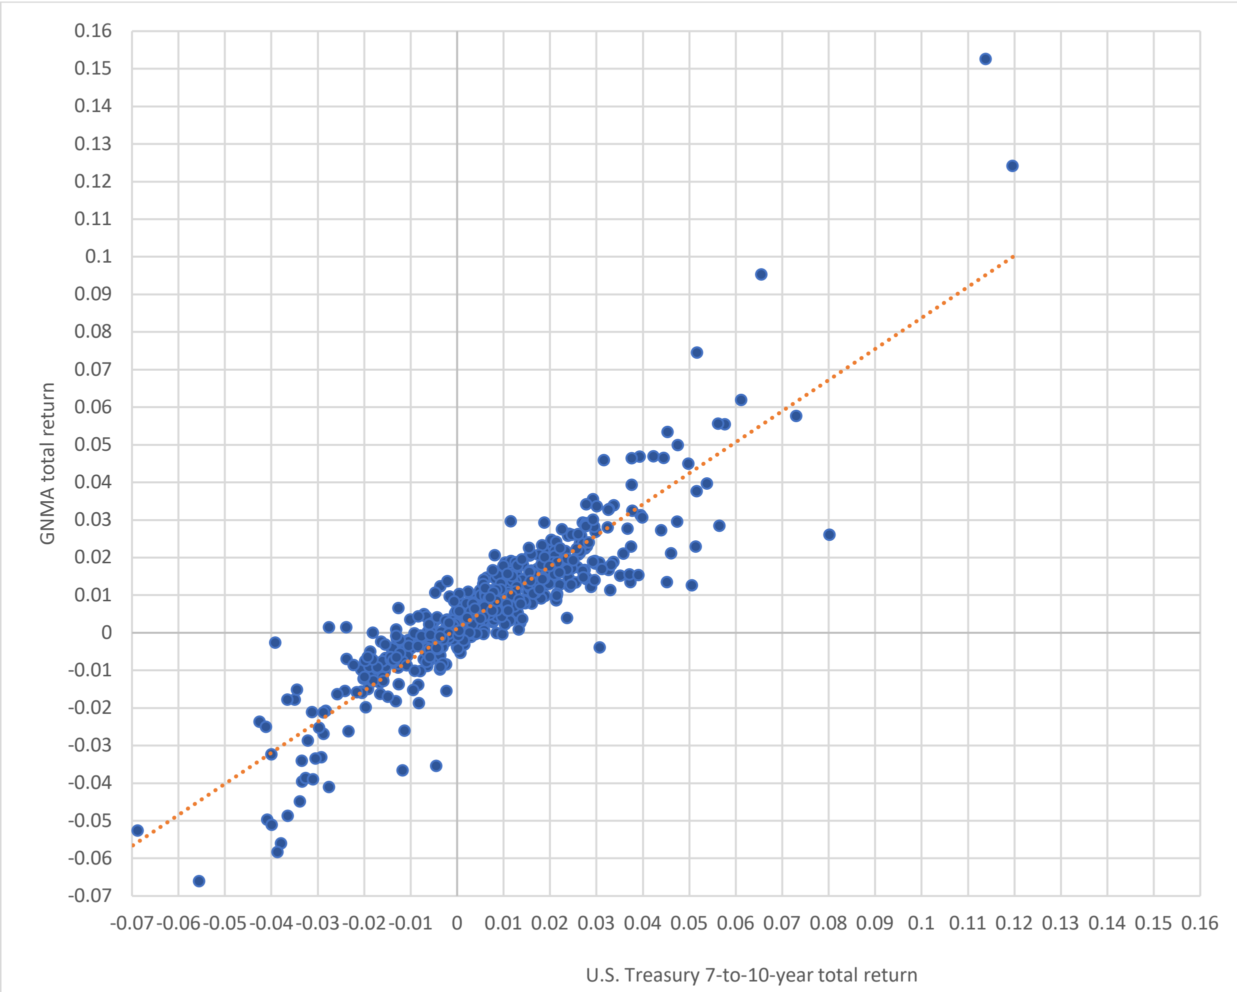

Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.
Note: GNMA = MBSs issued by the Government National Mortgage Association.

It can be difficult to assess the quality of a simple MBS total-return model. A high correlation of MBS returns with Treasury returns is important, and some evidence of negative convexity when interest rates fall is essential.

5.4 Summary

In this chapter we discussed the essential features of a good ESG. An ESG is a complex system and one that must evolve in response to changes in market fundamentals and regulatory requirements. A good ESG has some general characteristics that include the following:

- A good ESG has a sound foundation for the way the models are built and the way the variables are interrelated. It has a full range of modeled financial variables and multi-economy capability.
- A good ESG is capable of accommodating many types of calibration views across a wide variety of benchmarks. A good ESG produces simulation results that reflect a relevant view, i.e., one that is consistent with historical facts.
- A good ESG produces some extreme but plausible outcomes, which encapsulate historical behavior but do not go too far from market norms.
- A good ESG embeds realistic market dynamics. This requires agreement on a selection of stylized facts and institutional details.
- A good ESG is computationally efficient and numerically stable.
- A good ESG can meet the requirements of regulators and auditing firms.
- A good ESG has fast and robust recalibration capabilities.

Statistical criteria are also important in assessing the quality of an ESG. Statistical calibration criteria are usually numerically specified but can also be qualitative in nature. Statistical criteria belong to one of two broad categories:

- **Qualitative features:** An important first step in validation is to check that the most important qualitative stylized facts are satisfied by the simulated output.
- **Quantitative measures:** Tabular calibration targets will usually include targets for average levels and volatilities.

A path represents one possible future evolution of the economy and therefore represents one possible complete future “economic experience.” The importance of pathwise model behavior is that it is the simulated path that represents the way an insurance company will experience the evolution of the economy. If the overall distribution of returns for an asset class is correct but the pathwise behavior does not correspond to the nature of the fluctuations that we see in the historical record, then there is an issue with the model.

The fundamental process for real-world validation involves comparing calibration criteria against simulated model performance. The criteria used are both qualitative and quantitative. The chapter provides examples of both kinds of criteria applied to several kinds of situations.

References

Brigo, D., and F. Mercurio, *Interest Rate Models—Theory and Practice: With Smile, Inflation, and Credit*, 2nd ed., Berlin: Springer, 2007.

Chapter 6: Stochastic Processes and Dynamics for ESG Modeling

In this chapter, we explore methods for modeling certain economic and financial variables. While there are numerous possible approaches to modeling any particular variable, the methods described here are meant to be illustrative and relatively simple. They are, nevertheless, solid examples of modeling approaches and are representative of techniques employed in at least some ESGs. Thus, the chapter can provide a possible technical starting point for those wishing to develop their own ESG.

Much of the discussion here is enhanced and illustrated by reference to the financial scenario generator (an alternative name for an ESG) created by Ahlgrim, D'Arcy, and Gorvett (2004, 2005), supported by a research grant from the Casualty Actuarial Society and the Society of Actuaries. These references are meant to be illustrative only as one example of modeling in an ESG—an ESG that is publicly available and downloadable,[8] and that thus can be examined more closely and even utilized directly by readers if they are so inclined. This model will herein be referred to as the ADG economic scenario generator.[9]

In Section 6.1, we start by presenting some foundational mathematics associated with stochastic processes. Discrete- and continuous-time stochastic processes provide a basis for modeling the dynamics and the evolution over time of many economic and financial series. Section 6.2 identifies some of the econometric and statistical techniques useful for economic and financial analysis and modeling. Section 6.3 describes modeling approaches for certain key variables in an ESG, particularly interest rates and equity returns. Section 6.4 identifies several important considerations associated with modeling and using an ESG.

6.1 Stochastic Processes

Economic and financial variables are inherently stochastic: the future values they take on are uncertain prior to their actual emergence. Thus, the mathematical area of "stochastic processes" provides the basis for modeling the dynamics underlying these variables. In this section, we discuss the basics of stochastic processes, from simple discrete random variables to continuous-time processes.[10] The mathematics in this section provides the foundation for the illustrative models introduced in Section 6.3.

6.1.1 Discrete-Time Stochastic Processes

The discrete-time framework involves values of variables only at certain points in time. We start with a very simple type of discrete-time process.

---

[8] Available from https://www.casact.org/research/econ/.

[9] We also employ the ADG model in Chapter 7 as the basis for an illustrative discussion of model logistics such as input and output.

[10] A list of textbooks on stochastic processes and stochastic calculus can be found in Chapter 11.

Binomial Process

A binomial process is one whose value, over a fixed interval of time, can change by only one of two possible amounts. For example, a share of stock whose value is currently 100 may be modeled as either going up or down by 10% over the next year (to 110 or 90, respectively). The key parameter associated with such a binomial model is the probability of the stock going up versus going down in value—and the probability of a down movement would be equal to one minus the probability of an up movement. Note that while it is common to think of a binomial framework as having the potential for an "up" and a "down" movement, this need not be the case. The important point is that the process will branch to one of two future values; theoretically, those two values could both be up or both be down. Nevertheless, we often refer to "up" and "down" movements as meaning, respectively, a movement to the higher or lower future value.

Wiener Process

We can now enhance the binomial process framework described above by shortening the length of time over which the up-or-down movement occurs and compounding the binomial branching upon itself over consecutive time intervals. Thus, for example, instead of one up-or-down movement per year, there is one every six months, or one every three months, or one every Δt of a year—with a new branching occurring from each node during the next Δt time interval.

This describes a Markov stochastic process that is known as a *Wiener process*, or a simple *Brownian motion*. The idea is that, over a small interval of time, the process z changes slightly, according to the following properties:

$$\text{For } \Delta z = z(t + \Delta t) - z(t),$$

$$(i)\ \Delta z = \varepsilon\sqrt{\Delta t}, \ \ \text{where } \varepsilon \sim N(0,1)$$

(ii) Δz values are independent between different intervals of time Δt.

Since $\varepsilon \sim N(0,1)$ with $E[\varepsilon] = 0$ and $Var(\varepsilon) = 1$, the properties of Δz are

$$E[\Delta z] = E\left[\varepsilon\sqrt{\Delta t}\right] = E[\varepsilon]\sqrt{\Delta t} = 0,$$

$$Var(\Delta z) = Var\left(\varepsilon\sqrt{\Delta t}\right) = Var(\varepsilon)\left(\sqrt{\Delta t}\right)^2 = \Delta t.$$

We can also see that $z(t)$ is a martingale:

$$E[z(t + \Delta t) \mid z(t)] = z(t).$$

Now, if we let $z(t)$ be a stochastic process from t = 0 to T, we can use this structure to say something about the nature of the value distribution of the process at T, i.e., $z(T)$, given the value $z(t)$ at time 0. Let our time period of length T be made up of a large number n of subperiods of time length Δt—that is, let $\{0 = t_0, t_1, t_0\ , \ldots, t_n = T\ \}$ be a set of time-points evenly spaced between time 0 and T and $t_i - t_{i-1} = \Delta t$ for all $i = 1, \ldots, n$. Then,

$$z(T) - z(0) = \sum_{i=1}^{n} \left(z(t_i) - z(t_{i-1})\right) = \sum_{i=1}^{n} \varepsilon_i \sqrt{\Delta t} \quad \text{where } \varepsilon_i \sim iid\ N(0,1),$$

$$E[z(T) - z(0)] = \sum_{i=1}^{n} \sqrt{\Delta t}\, E[\varepsilon_i] = 0,$$

$$Var[z(T) - z(0)] = \sum_{i=1}^{n} \left(\sqrt{\Delta t}\right)^2 Var[\varepsilon_i] = n \cdot \Delta t = T,$$

$$z(T) - z(0) \sim N(0, T),$$

$$z(T) \sim N(z(0), T).$$

This last line is the "payoff" of this mathematical development of the Wiener process. It tells us that the change in value of a stochastic or random variable, looked at as a simple Brownian motion, has a normal distribution. More specifically, the volatility of the process is directly related to the length of time over which the process is evolved. This gives us a basis on which to analyze and predict economic and financial variables as well as a foundation for understanding and parameterizing the governing dynamics of those variables.

6.1.2 Continuous-Time Stochastic Processes

In the economic and financial worlds, data are gathered as discrete, rather than continuous, values: we accumulate data on processes by taking "snapshots" of the values of those processes at various moments in time. Thus, as addressed in Section 6.2, the tasks of parameterizing and calibrating models of economic processes generally involve looking at those processes in a discrete-time framework—since the data that will allow us to do those tasks are themselves discrete.

Nevertheless, there is often value in describing economic and financial variables as continuous-time processes. Conceptually, those variables do potentially change values and evolve continuously, and—importantly—the mathematics is often easier to work with on a continuous-time basis as opposed to a discrete-time basis. Indeed, developing a continuous-time mathematical framework can sometimes lead to closed-form solutions for describing future values. Thus, it is common in ESG development to use a continuous-time framework to describe a variable's underlying dynamics, but to recognize that it is the discrete-time analogue of that continuous-time specification that is actually employed when parameterizing and calibrating the model with actual data.

A continuous-time stochastic process z can be looked at as the limiting value of a discrete-time process, as the interval of time Δt approaches zero:

$$dz_t = \lim_{\Delta t \to 0} \{ z(t + \Delta t) - z(t) \}.$$

Basically, as $\Delta t \to 0$, we have $\Delta z \to dz$.

Now, we can describe a process by a continuous-time differential equation using standard notation from calculus.

Often, a stochastic process is modeled as having two parts: a deterministic term, and a stochastic or volatility term. For example, a generalized Brownian motion process x can be described mathematically as

$$dx = \alpha dt + \beta dz.$$

This process is often called the arithmetic Brownian motion (ABM). In this construction, the change in the process x depends upon a deterministic component (the first term on the right), which tends to push the value of x up or down according to the parameter α (called the "drift" of the process) and upon a volatility term (the second term on the right). In the volatility term, dz is a standard Brownian motion or Wiener process (it can be thought of as a random sampling from a standard normal distribution—recall the description of the Wiener process above), and β is the volatility parameter that indicates the relative size of random movements in the process. As mentioned above, it is important to keep in mind that this continuous-time construction, which is mathematically convenient, has a discrete-time analogue that can be employed when parameterizing or calibrating the process with data:

$$\Delta x = x(t + \Delta t) - x(t) = \alpha \Delta t + \beta \Delta z.$$

As with the simple Wiener process, this more general process can be characterized according to the mean and variance of its future values:

$$E[\Delta x] = E[\alpha \Delta t + \beta \Delta z] = \alpha \Delta t,$$

$$Var(\Delta x) = V(\alpha \Delta t + \beta \Delta z) = Var\left(\beta \varepsilon \sqrt{\Delta t}\right) = \beta^2 \Delta t,$$

$$\Delta x \sim N(\alpha \Delta t, \beta^2 \Delta t),$$

$$x(t + \Delta t) \sim N(x(t) + \alpha \Delta t, \beta^2 \Delta t).$$

Another important version of a Brownian motion is the geometric Brownian motion (GBM).

A GBM process S is mathematically described as[11]

$$dS = \mu\, S\, dt + \sigma\, S\, dz \text{ or } {}^{dS}/_{S} = \mu\, dt + \sigma\, dz.$$

For GBM, the discrete-time analogue is

$$\Delta S = \mu\, S\, \Delta t + \sigma\, S\, \Delta z \text{ or } {}^{\Delta S}/_{S} = \mu\, \Delta t + \sigma\, \Delta z.$$

[11] "S" is used here because stock processes are often modeled by geometric Brownian motion.

(Note that we have used the same μ and σ coefficients for both the discrete-time and continuous-time expressions. In reality, the values of both parameters would most likely be somewhat different between the two frameworks.)

The above stochastic differential equation for a GBM can be solved analytically:

$$S_t = S_0 \exp\left((\mu - \frac{\sigma^2}{2})t + \sigma z_t \right),$$

where S_0 is the value of the process at time zero and is regarded as a non-stochastic constant. Given S_0, S_t is said to follow a log-normal distribution with mean and variance given by

$$E(S_t) = S_0 e^{\mu t},$$

$$V(S_t) = S_0^2 e^{2\mu t}\left(e^{\sigma^2 t} - 1\right).$$

From the solution, we can write

$$dlog(S_t) = \frac{dS_t}{S_t} = \left(\mu - \frac{\sigma^2}{2}\right) dt + \sigma dz_t.$$

With a discrete-time version of this relation, we can identify the nature of the distribution of "return" of the process:

$$\frac{\Delta S}{S} \sim N\left(\left(\mu - \frac{\sigma^2}{2}\right)\Delta t, \sigma^2 \Delta t\right).$$

GBM is frequently used to model stock prices such as in the Black–Scholes model.

6.1.3 Other Dynamics

One other process that can be useful in modeling economic and financial stochastic processes is a jump process. Such a process can be incorporated into a Brownian motion process to create a mixed jump diffusion process that allows for a stochastic jump at random times. A common approach is to let λ be the jump rate (expected number of jumps per time period), let k be the expected jump size (as a percentage of the value of the underlying variable), that is, $E[y] = k$, and assume that the occurrence of jumps follows a Poisson process.

In such a construction, the frequency of jumps during the time interval Δt is equal to $\lambda \, \Delta t$. Then, the expected rate by which the underlying process changes due to jumps in a year is λk. For example, if stock prices are otherwise modeled with GBM, the addition of a jump process would result in the following diffusion process:

$$dS/_S = (\mu - \lambda k)dt + \sigma \, dz + y \, dN.$$

In this way, the overall mean drift of the process after the introduction of the jump process is the same as before jumps were added.

6.2 Econometric and Statistical Techniques

The previous section showed some common stochastic processes that can be used as a basis for modeling economic and financial variables. This type of direct modeling of one or more variables generally involves using a stochastic process framework. In this section, we briefly describe several econometric and statistical techniques that can be of value for at least two purposes: (1) modeling the relationships between variables, and (2) parameterizing and calibrating the ESG model. These two purposes mirror the two critical questions in economic modeling: the adequacy of the model in describing economic processes and the accuracy of the parameters used in the model.

6.2.1 Variable Relationships and Correlations

At some point, ESG developers may feel that one or more essential variables have been modeled—for example, nominal interest rates, real interest rates, or inflation—and that all other variables can be modeled effectively by linking to, or cascading from, those one or more key variables. Here are some possibilities for achieving such simulated variable relationships:

- **Direct linkages:** We can model the relationship between future values of different variables by simply making one variable a function of another. If perfect, lockstep correlation between the two variables is not desired (as it generally would not be), some type of error or volatility term could be included. This could take the form of a simple or multivariate regression equation with a random error term.

- **Correlation of volatility terms:** If two variables are modeled by different stochastic processes, a correlation factor between the two Wiener processes that provide volatility to each variable—i.e., the dz terms—can be incorporated into the simulation procedure. When more than two variables need to be correlated, a Cholesky decomposition of a correlation matrix can be used to generate correlated Wiener processes.

- **Time-varying correlations:** While obviously increasing the modeling sophistication and difficulty level, it is possible that correlations between variables can change over time. If historical evidence is found to that effect, then it may be desirable to include in the ESG a procedure to change correlation levels as simulations push further out into the future.

6.2.2 Parametrization, Calibration, and Understanding Interrelationships

Whatever modeling approach one uses for the various economic and financial variables in an ESG, the simulations emerging from the ESG will be only as good as the parameterization and calibration of the model derived from the data. In addition, a good model should help the user better understand the interrelationships between variables. These tasks can be supported by a variety of econometric and statistical techniques and considerations,[12] including the following:

[12] For a more detailed discussion about estimation techniques, see Chapter 4 and the references therein.

- **Maximum likelihood estimation**, which finds the parameter values that provide the highest probability of producing the historical data actually observed
- **Kalman filtering**, which recursively estimates parameters from a sequence of measurements that include noise
- **Leading and lagging correlations**, which can provide additional insight into variable interrelationships
- **Vector autoregression (VAR),** an extension of univariate autoregression that involves the modeling of multivariate time series data. It is a multiple-equation system in which each variable acts as the dependent variable in an equation.
- **Markov chain Monte Carlo (MCMC),** which estimates both the parameters and the latent variables of a model, by first generating a Markov chain of the parameters and latent variables, and then using those samples for estimation with the Monte Carlo method

6.3 Modeling Approaches for ESG Variables: Interest Rates, Macroeconomic Variables, and Equity Returns

In this section, we describe some of the techniques and considerations involved in modeling certain key economic variables.

6.3.1 Interest Rates

As one decides upon a model of the term structure of interest rates to use in an ESG, several issues need to be determined. Two of the most important are the type of model to use and whether to use a one-factor model or a multiple-factor model.

<u>Type of Model</u>

Interest rate models are often categorized into two groups: equilibrium and no-arbitrage models. This distinction is technically correct, but in practice it can be slightly misleading. According to this bifurcation, an "equilibrium model" is one for which the current term structure of interest rates is an *output*; a "no-arbitrage model" is one for which the current term structure is an *input*. The complication is that often a no-arbitrage model can be created as a simple *extension* of an equilibrium model.

An example involves the Hull and White (1990) model. In that paper, they consider extensions of the original equilibrium models of Vasicek (1977) and Cox, Ingersoll, and Ross (1985). Specifically, they create no-arbitrage versions of these models by allowing for time-dependent (but still deterministic) values of the parameters.

To illustrate, consider the original Vasicek model, which is a one-factor (one source of uncertainty) model of the short rate r:

$$dr_t = a\,(b - r_t)dt + \sigma\, dz.$$

This model of the evolution over time of the short rate (i.e., the force of interest) is a mean-reverting (called an Ornstein–Uhlenbeck) process with a deterministic drift and a stochastic volatility term (since dz is a standard Brownian motion process). This stochastic differential equation has three parameters:

- speed of mean reversion (a),
- long-run mean to which the short-rate reverts over time (b), and
- volatility (σ).

In the original Vasicek specification, each of these three parameters is time-homogeneous, or constant over time, and thus is rather restrictive in its ability to model different term structures. In particular, there are not enough degrees of freedom here to calibrate this model using the entire actual term structure as an input.

Hull and White considered the possibility that these parameters may, in fact, be time dependent. While this time heterogeneity can apply to any of the three parameters, consider a model that allows only the long-run mean parameter b to change over time:

$$dr_t = a\,(b_t - r_t)dt + \sigma\, dz.$$

This Hull–White model of the short rate differs from the original Vasicek model only in the parameter b—but by permitting that parameter to change over time, the additional flexibility and degrees of freedom allow the model to be calibrated with the starting term structure as an *input*. Thus, Hull–White is a no-arbitrage model, or a no-arbitrage extension of the original equilibrium model.

As always, there is a trade-off between more model parameters and a potential overfitting problem. In fact, much of the Hull and White paper is contextually motivated by the different purposes for which interest rate models are employed by practitioners versus academics. For ESG purposes, the primary use(s) to which the ESG is directed will be a significant factor in choosing an interest rate model. A model that allows for exactly fitting the current, beginning term structure—i.e., a no-arbitrage model for which the beginning term structure is an input in the calibration process—would seem to be preferable for rather short-term or trading purposes. For longer-term simulations, either type of model may be appropriate, as long as the interest rates it produces are well behaved over the long period of time.

Hull and White also develop an extension of the Cox–Ingersoll–Ross (CIR) equilibrium model[13] in the same way.

[13] The CIR (1985) model is a general equilibrium model in which the short rate evolves over time according to a mean-reverting diffusion process:

$$dr_t = \kappa\,(\theta - r_t)dt + \sigma\sqrt{r}\, dz.$$

Number of Factors

The number of stochastic factors in a term structure model indicates how many "sources of uncertainty" there are. The model will simulate the evolution of interest rates over time, and the movements in the term structure are generated by the (one or more) stochastic factors. In a one-factor model, the single stochastic factor is generally the instantaneous short rate of interest itself. In a two-factor model, another stochastic factor (or source of uncertainty) is added. This second factor can take several forms. For example, for a mean-reverting interest rate, the second source of uncertainty may be the long-term mean to which the short rate tends to revert. In other words, while the short rate itself is stochastic and varies over time, the long-run mean is also stochastic and follows its own diffusion process. Another possibility for the second factor in a two-factor model might be volatility, which could be hypothesized to evolve stochastically over time. Another possibility is to have the short rate as a sum of two or more stochastic factors (usually no more than three).[14]

Note that introducing a second or third "source of uncertainty" is different than allowing a *deterministic* parameter to vary over time, as mentioned in the last subsection with respect to extensions of equilibrium models. Here, a second or third factor involves making a parameter stochastic, with its own volatility or dz term that allows it to take on a variety of possible values over time—with each simulation of the process resulting in a different realization for the value of the factor.

The choice of how many factors a term structure model should have is a decision involving trade-offs.[15] Including more stochastic factors increases the ability of the model to reflect a wide variety of actual interest rate movements. Indeed, with a one-factor term structure model, the dynamics of the entire yield curve are completely driven by the single source of uncertainty. Resulting yield curve movements are subsequently constrained: yields of all maturities are perfectly correlated to the one stochastic factor and the range of potential yield curves is limited.[16] This may be felt to be too limiting for a particular organization's ESG purposes.

On the other hand, fewer factors may make the model simpler, more parsimonious, and more tractable. There is also the "philosophic" issue of whether a model should be capable of reproducing all historical observations. While that capability in and of itself sounds attractive and reasonable, the cost of doing so, in terms of additional model complexity and thus less causal transparency, may be deemed too high.

Ahlgrim–D'Arcy–Gorvett (ADG) ESG

In the development of their economic scenario generator, ADG desired a model with closed-form solutions. Ultimately, they selected the two-factor Vasicek model of the term structure of real

[14] In this case, the single stochastic factors cannot be identified with observable rates, and in the parameter estimation a filtering technique like a Kalman filter is needed (see Bolder 2001).

[15] A principal component analysis on yield historical series shows that three factors explain approximately 98% of the sample variation (see Litterman and Scheinkman 1991).

[16] Ahlgrim, D'Arcy, and Gorvett (2004).

interest rates.[17] The two factors in this model are the short-term interest rate r and the long-term rate l to which the short-term rate tends to revert. Mathematically, the diffusion processes for these two factors are:

$$dr_t = \kappa_1(l_t - r_t)dt + \sigma_1 dB_1,$$

$$dl_t = \kappa_2(\mu - l_t)dt + \sigma_2 dB_2,$$

where μ is the value to which the long-run mean l reverts in the long run, and the dB terms are standard Wiener processes.

The dynamics of the term structure model in ADG were parameterized by applying historical data to the discrete-time version of these formulas.

6.3.2 Equity

An important characteristic of historical equity returns is that they exhibit fat tails—that is, the probability of extreme values well above and well below the mean is greater than would be implied by assuming the returns follow a normal (or Gaussian) distribution. Over the last several decades, a number of "unusual" equity market movements have reinforced the idea that equity returns exhibit fat tails. For some observers, this has been interpreted as new information, but the fat tails phenomenon has been known for many years—at least since Fama (1965).

Figure 6.1. S&P 500 daily returns

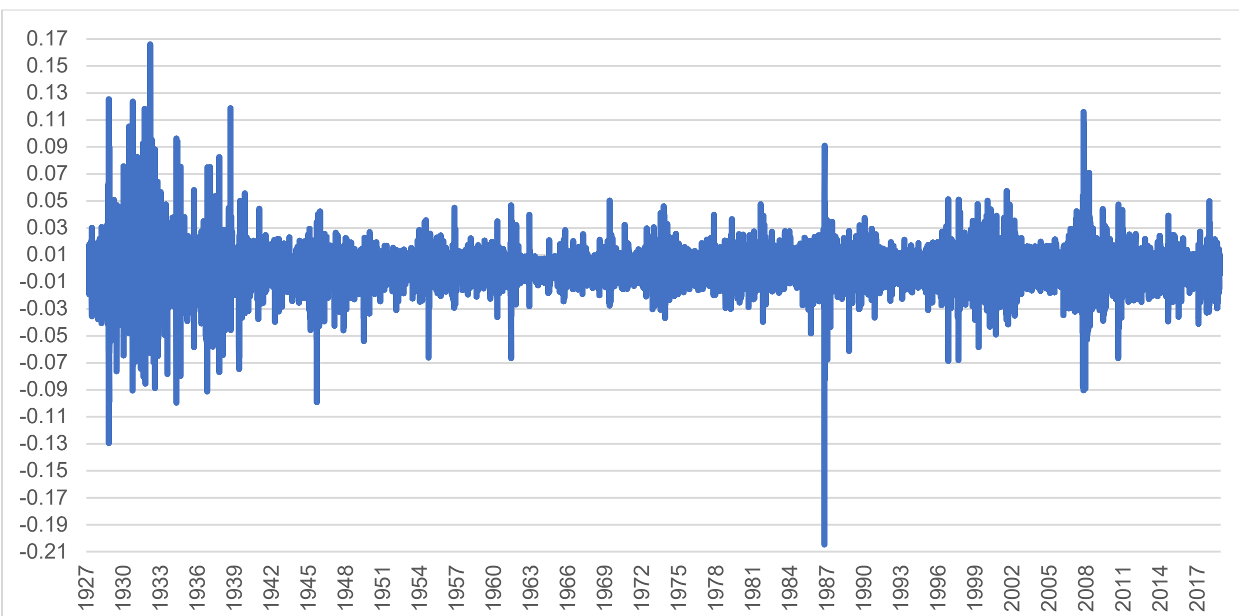

[17] ADG used a term structure model of real interest rates because they separately modeled inflation. When inflation and real interest rates were simulated, they were combined to produce nominal interest rates.

Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.

Another important characteristic of historical equity returns is that the volatility is not constant over time, but stochastic, and it shows clustering (see discussion in Chapters 4 and 5)—that is, periods of high volatility are followed by periods of lower volatility. See Figure 6.1.

A possible way to model fat tails and volatility clustering is by using a jump-diffusion model with stochastic volatility, like the Bates (2006) model:

$$\frac{dS_t}{S_t} = \left[\mu_0 + \mu_1 V_t - (\lambda_0 + \lambda_1 V_t)\bar{k}\right]dt + \sqrt{V_t}\left(\rho dW_{1t} + \sqrt{1-\rho^2}dW_{2t}\right) + (e^{\gamma_s} - 1)dN_t,$$

$$dV_t = [\alpha - \beta\, V_t]dt + \sigma\sqrt{V_t}dW_{1t},$$

where dS_t/S_t is the instantaneous asset return, V_t is the instantaneous variance, W_{1t} and W_{2t} are independent Wiener processes, $N_t$is a Poisson counter with intensity $\lambda_0 + \lambda_1 V_t$, γ_s is the random Gaussian jump, and $\bar{k}$ is the expected percentage jump size.

Another way to model the fat tails of equity returns is with a regime-switching model. ADG chose that method, based on the approach of Hardy (2001). The regime-switching model involves two regimes, high volatility and low volatility, where the likelihood of staying in the current regime or switching to the other regime is governed by transition probabilities, as in a Markov chain. With this approach, large short-term movements in the equity market can be attributed to a randomly selected value from the high-volatility regime. From this perspective, occasional large equity movements seem more reasonable, compared with treating the market as having a single regime and thus a level of volatility that is in some sense an "average" between a high- and low-volatility regime.

While equity returns can be modeled directly in the ADG model, the preceding description applies to the simulation of *excess* equity returns above a risk-free nominal interest rate. Full equity return rates were then simulated by combining the excess return with the simulated risk-free nominal interest rate (which in turn was produced by adding inflation to a simulated real interest rate).

6.3.3 Some Other Economic and Financial Variables

Inflation

One of the most visible and closely watched of economic variables, inflation can be, and has been, modeled in numerous ways—including some approaches that are really quite simple mathematically. For example, because inflation tends to exhibit persistence through time, a simple autoregressive process involving the current value and one or more lagged values has modeling appeal. Other types of regression specifications are also common modeling efforts for inflation.

The ADG model employs an Ornstein–Uhlenbeck mean-reverting process of the form

$$dq_t = \kappa(\mu_q - q_t)dt + \sigma dB_q,$$

where q is inflation, κ is the speed of mean reversion, $\mu_q$is the long-run rate toward which the process tends to revert, and σ is a volatility parameter. The above formula is a continuous-time version of the inflation process; the model is parameterized using historical inflation data and a discrete-time version of the formula. The above model and more in general continuous-time models are quite appealing because one can derive closed or semi-closed formulas for the term structure of expected inflation and inflation-linked bonds and derivatives.[18] But such kinds of models might not be able to capture a highly persistent autocorrelation over long time periods.

Equity Dividend Yields

Another example of using an autoregressive mean-reverting process to model the dynamics of a financial variable in ADG is with respect to equity dividend yields. Similar to the approaches used by Wilkie (1984) and Hibbert, Mowbray, and Turnbull (2001), the process used to model the log of the dividend yield y is

$$d(\ln y_t) = \kappa_y(\mu_y - \ln y_t)dt + \sigma_y dB_y.$$

Unemployment

A final economic variable worth mentioning from the ADG model is unemployment, which is modeled with a first-order autoregressive process. For unemployment, a common relationship with other economic variables is through the Phillips curve, which reflects an inverse relationship between unemployment and inflation. Even if the Phillips curve has not worked well recently, it is still a standard framework used in central banks.[19]

6.4 Important ESG Modeling Considerations

In this section, we briefly identify a number of important considerations for the builder and user of an ESG.

- **Data periods:** For many models, calibration and parameterization can be quite sensitive to the period of time from which historical data are collected and analyzed. Within this consideration are several sub-issues:
    - *Stability versus responsiveness*: As a common trade-off and concern in general actuarial work, it is important to consider where the happy medium is between a long period of data (enhancing stability) and a recent shorter data period (that promotes responsiveness to more recent conditions).
    - *Different points in the business cycle:* Over the course of several decades, the economy and financial markets can go through a variety of stages, both good and bad. Some models can be very sensitive to which data periods are selected as the

[18] See, for instance, Jarrow and Yildirim (2003).

[19] See "The Phillips Curve May Be Broken for Good," *The Economist,* November 1, 2017.

basis for parameterizing the model. A good guide is to try to use a long enough period of data to encompass several different economic environments, but one should still use a tremendous amount of judgment in making such decisions.

- **Philosophy:** To guide decisions such as the amount and length of data to use, it is helpful to have a firm understanding of the purpose and philosophy behind the creation and use of an ESG. Particularly important is to understand what the ESG is expected to "represent." Should it be capable of simulating the range and appearance of actual historical economic and financial data? Should it be merely representative of the general kinds of patterns that have emerged historically in the data? Does the user want the model to simulate economic conditions and values that, while reasonable, have never actually occurred in reality? Questions such as these can help guide one's development, interpretation, and evaluation of an ESG model.

- **Risk-neutral and real-world frameworks:** The risk-neutral and real-world frameworks coexist and have a place in the ESG environment. For example, the real-world price for a financial option is the expected value of the option's possible future values (or potential cash flows) under a general assumption that all investors are risk neutral—i.e., under a risk-neutral process. Meanwhile, certain regulatory authorities require a market-consistent (risk-neutral) framework for the valuation of insurance liabilities.[20]

- **Simulation versus scenario analysis:** Ideally, an ESG will have the capability of either simulating future economic values stochastically or producing a specific economic scenario. The latter capability is important in addressing stress-testing effects associated with economic scenarios promulgated by regulators.

- **Negative values:** For certain economic variables—such as nominal interest rates—a decision will need to be made regarding whether to allow negative values. This can sometimes be a contentious and difficult issue, with considerations involving historical precedents (or the lack thereof), economic rationality, and mathematical and modeling reality.

6.5 Summary

In this chapter, we explore ways of modeling certain economic and financial variables. For illustration, we make reference to the financial scenario generator created by Ahlgrim, D'Arcy, and Gorvett (ADG economic scenario generator).

First, we discuss the basics of stochastic processes, from simple discrete random variables to continuous-time processes. The mathematics in this section provides the foundation for the illustrative models introduced in Section 6.3.

[20] More details on real-world and risk-neutral frameworks can be found in Chapter 1.

The discrete-time framework involves values of variables only at certain points in time, but modelers also often describe economic and financial variables as continuous-time processes. A continuous-time framework can describe a variable's underlying dynamics, but a discrete-time analogue of that continuous-time specification is actually employed when parameterizing and calibrating the model. Finally, one other process that can be useful in modeling economic and financial stochastic processes is a jump process.

Modelers use several econometric and statistical techniques for at least two purposes: (1) to model relationships between variables, and (2) to parameterize and calibrate the ESG models. One or more key variables may be modeled with all other variables modeled effectively by linking to, or cascading from, those key variables. Some possibilities for achieving such simulated variable relationships include direct linkages, correlation of volatility terms, correlation matrices, and time-varying correlations. In addition, a good model should help the user achieve a better understanding of the interrelationships between variables. These tasks can be supported by a variety of econometric and statistical techniques and considerations—for example, maximum likelihood estimation, Kalman filtering, leading and lagging correlations, and Vector autoregression, which involves the modeling of multivariate time series data.

Certain techniques and model requirements are more appropriate for modeling specific economic variables. When deciding upon a model of the term structure of interest rates, two important considerations are the type of model to use (general equilibrium or arbitrage free) and whether to use a one-factor model or a multiple-factor model. The number of stochastic factors in a term structure model indicates how many "sources of uncertainty" there are. An important characteristic of historical equity returns is that they exhibit "fat tails"—that is, the probability of extreme values well above and well below the mean is greater than would be implied by assuming that the returns follow a normal (or Gaussian) distribution. Inflation is one of the most visible and closely watched of economic variables, and it can be, and has been, modeled in numerous ways—including some approaches that are really quite simple mathematically. For example, because inflation tends to exhibit persistence through time, a simple autoregressive process involving the current value and one or more lagged values has modeling appeal. Other types of regression specifications are also common in modeling inflation. Autoregressive processes may also be useful in modeling other variables, such as equity dividend yields and unemployment.

Finally, we briefly identify a number of important considerations for builders and users of an ESG: the sensitivity of the data period to stability and responsiveness; a clear understanding of the philosophy or purpose of the ESG; the choice of risk-neutral or real-world framework; the importance of simulation or scenario analysis; and whether negative values will be permitted in modeling certain economic variables.

References

Ahlgrim, K. C., S. P. D'Arcy, and R. W. Gorvett, "Report on Modeling of Economic Series Coordinated with Interest Rate Scenarios," Arlington, VA: Casualty Actuarial Society, 2004, https://www.casact.org/research/econ/.

Ahlgrim, K. C., S. P. D'Arcy, and R. W. Gorvett, "Modeling Financial Scenarios: A Framework for the Actuarial Profession," *Proceedings of the Casualty Actuarial Society* 92, 2005, pp. 177–238.

Bates, D. S., "Maximum Likelihood Estimation of Latent Affine Processes," *Review of Financial Studies* 19:3, 2006, pp. 909–965.

Bolder, D., "Affine Term-Structure Models: Theory and Implementation," Staff Working Paper 2001-15, Ottawa: Bank of Canada, 2001.

Cox, J. C., J. E. Ingersoll, and S. A. Ross, "A Theory of the Term Structure of Interest Rates," *Econometrica* 53:2, 1985, pp. 385–407.

Fama, E. F., "The Behavior of Stock-Market Prices," *Journal of Business* 38:1, 1965, pp. 34–105.

Hardy, M. R., "A Regime-Switching Model of Long-Term Stock Returns," *North American Actuarial Journal* 5:2, 2001, pp. 41–53.

Hibbert, J., P. Mowbray, and C. Turnbull, "A Stochastic Asset Model & Calibration for Long-Term Financial Planning Purposes," Technical Report, Barrie & Hibbert Limited, 2001.

Hull, J., and A. White, "Pricing Interest-Rate Derivative Securities," *Review of Financial Studies* 3:4, 1990, pp. 573–592.

Jarrow, R., and Y. Yildirim, "Pricing Treasury Inflation Protected Securities and Related Derivatives using an HJM Model," *Journal of Financial and Quantitative Analysis* 38:2, 2003, pp. 337–358.

Litterman, R., and J. A. Scheinkman, "Common Factors Affecting Bond Returns," *Journal of Fixed Income* 1:1, 1991, pp. 54–61.

Vasicek, O., "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5, 1977, pp. 177–188.

Wilkie, A. D., "A Stochastic Investment Model for Actuarial Use," *Transactions of the Faculty of Actuaries* 39, 1984, pp. 341–403.

Chapter 7: Illustrative Modeling of Three Key ESG Components

In this chapter, we discuss modeling issues associated with three key—and generally the most significant and impactful—ESG variables: inflation, interest rates, and equity returns. These variables are considered from the standpoint of modeling design and structure, as well as with respect to calibration, verification, and validation.

The jumping-off point for the examination of each of the three variables will be the publicly available ESG model of Ahlgrim, D'Arcy, and Gorvett (2004, 2005, 2008), hereafter referred to as the ADG model.[21] Undoubtedly other more sophisticated ESG models exist, but starting our discussion of each variable with how that variable is treated in the ADG model has several advantages:

- The ADG model stems from a joint Casualty Actuarial Society/Society of Actuaries request for proposal (CAS/SOA RFP, issued in 2001) and thus has from the beginning been intended for an actuarial/insurance audience.
- The model was, and continues to be, publicly available, as is its documentation.[22]
- The model is relatively simple compared with what we understand about some proprietary ESGs.
- The platform for the model is also simple and accessible: the model itself is built in Microsoft Excel and is run with the simulation software @RISK (a Palisade product), an Excel add-on.

While ADG acts as a starting point for our discussion of each variable, we will go beyond ADG in this chapter and address the following:

- Section 7.1: Overview of the general structure of the ADG model
- Section 7.2: Discussion of each of three key economic variables modeled by an ESG—inflation, interest rates, and equity returns. For each variable, the discussion covers three items:
    - ADG's approach to modeling the variable;

[21] The ADG model was used as the basis for several sample formulas for generating certain economic and financial variables in Chapter 6. Those samples were specific to that chapter's purpose, which was to describe how stochastic-process mathematics is used to model the dynamics of such a variable over time.
[22] The model and its documentation are available at https://www.casact.org/research/econ/.

- one or two illustrative issues associated with the parameterization and calibration of a model for that variable; and
- brief identification of some alternative modeling structures for the variable.

- Section 7.3: Discussion of considerations involved with verifying and validating a calibrated ESG model
- Section 7.4: Examples of sample model output
- Section 7.5: Some potential sources for economic and financial data

Before proceeding, let's define several important words that are sometimes confused for one another, even though they each refer to something different. For our purposes, we will employ these four terms according to the following definitions:

- *Parametrization*: identifying specific parameters (or structural characteristics) of a model. One does this through estimation from historical data, experimentation, or judgment.
- *Calibration*: adjusting modeling parameters in order to improve the model's agreement with actual current data—e.g., with the current term structure
- *Verification*: confirming that the implementation of the model is consistent with the designer's conceptual structure
- *Validation*: determining how accurately the model represents the real world, given the model's purpose

7.1 Overview of the ADG model

To better appreciate the ESG model context of the three key economic variables—inflation, interest rates, and equity returns—on which we focus in Section 7.2, this section briefly describes the overall design and structure of the entire ADG model.

To set the stage, here are some quotations from the original model. They reflect the thinking that led to the RFP that initiated that research and the conceptual approach employed by the ADG model authors:

> The goal of this project is to provide actuaries a model for projecting economic indices with realistic interdependencies among the variables.[23]

[23] Ahlgrim, D'Arcy, and Gorvett (2004, sec. 2), from original CAS-SOA RFP.

> The model is intended to be a useful tool, general enough to pertain to a variety of actuarial applications including, but not limited to, dynamic financial analysis, cash flow testing, solvency testing, stress testing, reserving, and pricing.[24]
>
> A key part of dynamic financial analysis modeling is the reasonable representation of future economic indices, to model asset and liability risks. In cash flow testing, plausible future scenarios must be created to include or be consistent with plausible values of a variety of economic indices. A standardized approach to this problem would be an important step in providing guidance to practicing actuaries.[25]

When designing an ESG, certain high-level perspectives can both provide useful guidance for developers and help with understanding the model for users. One perspective is a schematic layout of the individual modeling components of the ESG: a chart that summarizes the sequential process for simulating the various economic and financial variables, and pictorially displays the various interrelationships and dependencies between those variables. We include such a schematic diagram of the ADG model in Appendix 7.A at the end of this chapter.

Another important perspective is how the overall model is structured from beginning to end, from user inputs through calculated outputs. Since the ADG model is built in Excel, the logic and algorithmic progression of the model is transparent and can be followed by paging through the various sheets of the workbook. As Ahlgrim, D'Arcy, and Gorvett (2004) describe in greater detail, the ADG model contains the following eight worksheets:

- *@RISK Correlation*: a correlation matrix specifying the interrelationships between the modeled economic variables
- *Scenarios*: allows users to specify particular scenarios for scenario and sensitivity testing
- *ModelInput*: user input regarding the selected parameter values associated with each modeled economic variable
- *StochProcs*: the calculations underlying the model's projections, with 50 years of simulated variable values being projected
- *OutputIntRates*: output and reporting of simulated real interest rates, nominal interest rates, and inflation
- *OtherOutput*: output and reporting on other variables, including large- and small-stock returns, dividend yields, real estate, and unemployment
- *IntRateChart* and *InitTermStructure*: two sheets, based on user inputs above, that exhibit (graphically and numerically, respectively) the implied starting term structure

[24] Ahlgrim, D'Arcy, and Gorvett (2004, app. A).

[25] Ahlgrim, D'Arcy, and Gorvett (2004, sec. 2), from original CAS-SOA RFP.

Finally, as a last piece of background on the ADG model, Appendix 7.B at the end of this chapter summarizes the key stochastic process and regression equations for each of the modeled economic variables. As with all discussion of the ADG model in this chapter, much more detail is available from the various ADG publications.

7.2 Modeling and Calibrating/Parameterizing the Key Economic Variables

In this section, we consider issues associated with how each of the three selected key economic variables is modeled and parameterized/calibrated. We start with how the variable is modeled in the public-access ADG model, and then extend the discussion beyond that particular model example.

Using the ADG model as a starting point presents an interesting opportunity. When it was initially developed and published in the early 2000s, the ADG model was parameterized based on data available at that time. In the nearly two decades since then, a lot has happened—socially, culturally, politically … and economically. This presents us with an opportunity to emphasize that, when using an ESG model, it is incumbent upon the user to determine whether any earlier parameterization and calibration underlying the model remains adequate, or whether it needs to be recalibrated. Almost certainly the latter approach is both appropriate and necessary—especially if one is interested in a model that seems to "start" from current economic reality.

For each of the three key variables, we look at

- the modeling structure used in the ADG model;
- issues of calibration/parameterization; and
- alternative modeling structures.

7.2.1 Inflation

Model Structure

As Chapter 6 documents, the ADG model uses an Ornstein–Uhlenbeck mean-reverting process to simulate the dynamics of inflation:

$$dq_t = \kappa_q(\mu_q - q_t)dt + \sigma_q dB_q.$$

This is the continuous-time version of the stochastic process representation for the evolution of inflation (q) over time. Since both historical inflation data and the simulated inflation series from the model are discrete, this continuous-time expression must be converted to a discrete-time version:[26]

$$\Delta q_t = q_{t+1} - q_t = \kappa_q(\mu_q - q_t)\Delta t + \varepsilon_q\sigma_q\sqrt{\Delta t}.$$

[26] Following Ahlgrim, D'Arcy, and Gorvett (2004, sec. 5). See that document for additional details.

Thus,

$$q_{t+1} = q_t + \kappa_q(\mu_q - q_t)\Delta t + \varepsilon_q \sigma_q \sqrt{\Delta t}$$

$$= \kappa_q \Delta t \cdot \mu_q + (1 - \kappa_q \Delta t) q_t + \varepsilon_q \sigma_q \sqrt{\Delta t}.$$

This equation is then the basis for the following regression:

$$q_{t+1} = \alpha + \beta \cdot q_t + \varepsilon'_{qt},$$

where $\alpha = \kappa_q \Delta t \cdot \mu_q$ and $\beta = (1 - \kappa_q \Delta t)$. This allows us to calibrate the mean reversion parameters based on the results of the regression:

$$\kappa_q = \frac{1 - \beta}{\Delta t}$$

and

$$\mu_q = \frac{\alpha}{1 - \beta}.$$

Ultimately, the ADG model uses this structure to develop a "term structure" of inflation that reflects expected inflation rates over various time horizons, using the Vasicek (1977) formula for the time *t* price of a bond that matures at time *T*:

$$P^q(t, T) = A(t, T) e^{-r\,B(t,T)},$$

where *A* and *B* are functions of the parameters of the inflation process.

Parametrization and Calibration

The first two decades of the 21st century have represented a rather different epoch than previous decades, in terms of the values and volatilities of financial and economic variables. This has been emphatically the case for inflation, which has progressed from historically high levels in the late 1970s and early 1980s to sustained low levels in the 2010s. This has a great many implications, but we mention two items in particular. First, there is the danger of believing that the future will always mimic the present. That may end up being true—especially in the short term—but the designer and user of an ESG must recognize (not just conceptually, but in practice) that such a present-sustaining future is but one of an infinitude of possibilities.

Second, and this is certainly related to the first, when data include major changes in a variable over time, the indicated parameters of the resulting model will depend significantly on the time period of the data to which the model is calibrated. While using a longer period of data will cover a wider variety of economic conditions, it can also give unwarranted weight to older conditions that may not be as relevant for the (at least short-term) future.

To demonstrate parameterization here, as was done in ADG (2004), we used the equation structure above and ran regressions against annual inflation rates, as calculated by

$$ln\left(\frac{CPI_t}{CPI_{t-1}}\right).$$

Not surprisingly, the parameter estimates for the inflation process vary greatly depending upon the data period used in the regression (Table 7.1).

Table 7.1. Inflation—parameter estimates based upon regressions

Data Period	κ_q	μ_q	σ_q
1968–2018	0.250	0.038	0.019
1978–2018	0.259	0.028	0.016
1988–2018	0.660	0.024	0.011
1998–2018	1.069	0.021	0.010

While there may be no "objectively correct" data period for parameterization and calibration, it is important to keep in mind the purpose of the ESG in a particular application. For example, if an ESG is being deployed in support of short-run operational decision-making, a shorter and more recent data period may better reflect the dynamics of economic and financial variables for that purpose. If long-run strategic planning is the ESG's primary focus, then a longer data period may provide a better range of variable values over a time period that may include several economic or business cycles.

Alternative Structures

Over the decades, models of inflation have been frequent and numerous, and have involved a wide variety of techniques. Historically, inflation has been characterized by persistent autocorrelation, and that is a key point of validation for modeling efforts. Also, because of the sensitivity of many types of insurance losses to inflation, the persistence of inflation has the potential to significantly affect losses over the long term.

In addition to the mean-reverting model of inflation used in ADG, the following is a sample of other time series structures analysts have commonly applied to inflation:

- Autoregressive (AR(p)) process: $X_t = \beta_0 + \beta_1 X_{t-1} + \cdots + \beta_n X_{t-n} + \varepsilon_t$, where $\varepsilon_t \sim N(0, \sigma^2)$ and $E[\varepsilon_t, \varepsilon_s] = 0$. This structure directly reflects the stylized fact of autocorrelation in the historical inflation process.

- Moving average (MA(q)) process: $X_t = \mu + \alpha_0 \varepsilon_t + \alpha_1 \varepsilon_{t-1} + \cdots + \alpha_q \varepsilon_{t-q}$, where $\varepsilon_t \sim N(0, \sigma^2)$ and $E[\varepsilon_t, \varepsilon_s] = 0$. This structure also addresses the persistence of inflation, through a linear regression of current and past noise or shocks to the process.

- Autoregressive moving average (ARMA(p,q)) process: $X_t = \beta_0 + \beta_1 X_{t-1} + \cdots + \beta_n X_{t-n} + \varepsilon_t - \alpha_0 \varepsilon_t - \alpha_1 \varepsilon_{t-1} - \cdots - \alpha_q \varepsilon_{t-q}$.

- Autoregressive integrated moving average (ARIMA) process: here, "integrated" reflects the fact that a nonstationary time series has been "differenced" to make it stationary.

7.2.2 Real Interest Rates

Model Structure

As mentioned in Section 6.3, ADG employed a double-mean-reverting two-factor Vasicek model for the term structure of real interest rates (which is a simple form of the Hull–White [1990] model):

$$dr_t = \kappa_1(l_t - r_t)dt + \sigma_1 dB_1,$$

$$dl_t = \kappa_2(\mu - l_t)dt + \sigma_2 dB_2,$$

where r is the short-term real interest rate, which tends to revert to the long-run mean l, μ is the value to which the long-run mean l tends to revert, and the two dB terms are standard Wiener processes. Nominal rates were then determined by combining projection of the real interest rates with those for inflation.

Several reasons led to ADG's adopting this two-factor model, among them:

- One-factor models generally result in 100% correlation between the interest rates of the term structure. Thus, the same shape of the yield curve will attend any two instances of the same simulated short rate. A two-factor model provides for more flexibility in the simulated rates and term structures.
- A model that provides for closed-form solutions will allow a user to track the points on the yield curve during a simulation. This is useful for certain applications that may be related to one or more specific yield curve maturities.

Ultimately, ADG use this process (as do Hibbert, Mowbray, and Turnbull 2001) to derive closed-form solutions for bond prices:

$$P^r(t,T) = A^r(t,T)e^{-r_t B_1(t,T) - l_t B_2(t,T)},$$

where A, B_1, and B_2 are functions of the parameters of the real interest rate process.

Then, to model nominal interest rates (i), the term structures of inflation and the real interest rate are combined:

$$P^i(t,T) = P^r(t,T) \times P^q(t,T).$$

Parameterization and Calibration

The dynamics of the real interest rate term structure model in ADG were parameterized by applying historical data to the discrete-time version of these formulas, involving regressions on monthly data. However, two noteworthy issues had to be dealt with:

- Since real interest rates are not directly observable in the market, they were determined on an ex post basis, as the difference between observed nominal interest rates and annualized monthly inflation.[27]

- Since the regressions were based on two formulas, one for the short rate and one for the long rate, and because one (the short rate) was dependent on the other (the long rate), ordinary least squares regression could not be used. Instead, ADG employed a two-stage least squares approach.

More generally, the specifics of calibrating an interest rate model depend largely on the particular model employed and its purpose. For example, if one intends to use a model for trading and pricing derivatives, prices of those types of instruments should be an input into the calibration process. If a two-factor model is used, involving the short rate and volatility as the two factors, data associated with both—such as a panel of relevant derivatives prices or a historical series of the yield curve—should both be inputs into the calibration process.

Alternative Structures

Practitioners have proposed many structures, involving many dimensions, for the modeling of interest rates and term structures, such as, for example,

- equilibrium versus no-arbitrage models;
- one- versus multiple-factor models;
- short-rate versus Heath–Jarrow–Morton (1992) framework model versus LIBOR market model; and
- modeling nominal versus real interest rates.

The selection of a particular interest rate model to use in an ESG is a function of many considerations:

- the trade-off between model flexibility and simplicity of use and parameter estimation;
- well-behaved and realistic simulations of future rates;
- the ability to observe relevant data necessary for parameterization; and
- sufficiently good fit to the data.

[27] In some economies, inflation-linked bonds are actively issued and traded, and they can be used to extract a term structure of real rates (see, for instance, Christensen, Dion, and Reid [2004] or Gürkaynak, Sack, and Wright [2008]).

Once the modeler has narrowed these preferences down, there exists an extensive literature on interest models, their use, and their estimation.

7.2.3 Equity Returns

Model Structure

As mentioned in Section 6.3, ADG use a regime-switching process to model equity returns. This regime-switching model, based on the approach of Hardy (2001), involves two regimes, high volatility and low volatility, where the likelihood of staying in the current regime or switching to the other regime is governed by transition probabilities, as in a Markov chain. With this approach, large short-term movements up or down in the equity market are the result of a randomly selected value from the high-volatility regime—as opposed to a single regime from which both common small movements and occasional large movements emerge.

An important attribute of regime-switching models is that they are consistent with fat tails—i.e., probabilities of extreme values well above and well below the mean being greater than if returns followed a normal (or Gaussian) distribution. This is one of the characteristics of the real-world history of equity returns. Indeed, the fat tails phenomenon for equity returns has been known for many years—at least since Fama (1965).

In ADG, the regime-switching equity returns model is used for simulating *excess* equity returns above a risk-free nominal interest rate. Full equity return rates are then simulated by combining the excess return with the simulated risk-free nominal interest rate (which in turn was produced by adding inflation to a simulated real interest rate).

Parameterization and Calibration

To determine parameters for the equity regime-switching model, ADG modeled the large-stock and small-stock return-generating processes as separate and independent, and used the approach of Hardy (2001) to maximize the implied likelihood function. This involves assuming that the stock prices in each regime are independent and log-normally distributed, and thus the log-returns $ln(S_t/S_{t-1})$ have independent and identical normal distributions. This is saying that stock prices follow geometric Brownian motion, and thus is consistent with the Black–Scholes model, which does well for short-term variations in equity returns. However, Hardy, Freeland, and Till demonstrate that "over longer terms, the [independent log-normal model] is generally rather thinner tailed than the data" (2006, pp. 80). This has implications for both model selection and calibration.

Although this *User's Guide* is concerned with property/casualty risks and the use of ESGs to help evaluate and manage such risks, one should not ignore the guidelines in place for the calibration of equity models in the context of life insurance and annuities. Organizations such as the American Academy of Actuaries (AAA) and the Canadian Institute of Actuaries (CIA) have promulgated acceptable models and calibration criteria that equity models should meet. For example, the CIA (2017) mentions four models for equity calibration that "remain appropriate":

- Log-normal model
- Two-factor regime-switching log-normal model
- Two-factor regime-switching drawdown model
- Correlated stochastic volatility log-normal model

The AAA promulgates numerical criteria that models must meet, relating to historical equity return characteristics such as negative skew, volatility clustering, and fat tails.

Alternative Structures

As with inflation and interest rates, many models have been proposed and used to represent equity returns. Common approaches might involve positing a geometric Brownian process (as underlies the Black–Scholes model), regime switching, and stochastic volatility models with jumps. Also, as alluded to in the prior subsection, the selection of an equity model may require consideration of requirements established by regulators or industry authorities.

In general, examples of some of the more common models of equity prices and/or returns (or of the equity risk premium, which is that part of the equity return that is in excess of the risk-free rate) might include, beyond those mentioned above, the following categories:

- Rolling n-period averages of historical stock index returns
- A dividend discount model: $P_t = \sum_{t=1}^{\infty} {E[D_t]}/{(1+i)^t}$, where D is the dividend payment. When assuming that the dividend increases by a constant percentage each year, the formula for the price is simply that for a geometrically increasing perpetuity. There are a number of suggested models that involve different ways of estimating earnings and/or the growth rates of dividends.
- Time-series regressions, on a variety of current or lagged explanatory variables
- Cross-sectional regressions, on various economic indicators and risk factors

7.3 Verifying and Validating a Calibrated ESG Model

Elsewhere in this *User's Guide*—particularly in Chapters 4 and 5—the validation process is described and discussed. In this chapter, we provide some brief illustrative examples of possible retrospective validation efforts, and some general guidance on the verification and validation processes.

Before proceeding, we emphasize a few important items with respect to the validation process:

- **Visualization:** Whereas quantitative measures of validation and goodness of fit can be extremely useful, a visual approach to reasonableness also adds great value. For example, comparing histograms of simulated output with actual economic variable emergence can help provide a strong intuitive sense of how reasonable a model's parameterization is. Since the human mind can be quite skilled at recognizing patterns, visual comparisons can be a good supplement to goodness-of-fit statistics. We make use of visualization in this section, with respect to validation of both inflation and equity returns.

- **Purpose of the ESG:** As is true with so many aspects of analytical work, the best approach to validation depends upon the reasons for using an ESG. The objectives of an ESG-related project—e.g., short- or long-term planning; emphasis on one functional area versus others—should help inform the selection of the most effective acceptance criteria.

- **Holistic:** Validation is a holistic process. It is useful, even important, to examine and validate individual modules within an ESG (as we illustrate in this section with respect to inflation and equity returns), but ultimately the entire model—including its variable interrelationships—should be assessed for reasonableness as an integrated whole.

- **Representative of history:** When using historical data to model and project future values, a philosophical question arises: do we want the model to in some sense "reproduce" history? While such a goal might be straightforward to validate—simply observing how well the simulation of future variables mimics and reflects historical patterns and relationships—that need not be the objective behind modeling. Rather, we may prefer that the simulated variable values be *representative of* history. While such an objective is challenging to define and validate, both quantitative and visual approaches can be taken to assess the model according to this goal.

7.3.1 An Illustrative Validation Test of Inflation

By examining the simulations of the future evolution of inflation, using the new parameters determined for the inflation variable described in Section 7.2, we can get an idea of how well our parameterized model represents history.

Recall that we developed new inflation parameters based on data through 2018. How do those new parameters perform when we simulate future inflation, say, 10 times for the next 20 years? And how do those simulations look relative to the history of inflation that went into the calibration and parameterization of the inflation module?

For illustrative purposes, we focus on a visual inspection of historical and simulated future inflation. For example, suppose we use the 1978–2018 data period parameter estimates. Figure 7.1 shows that history of annual inflation rates, plus ten 20-year projections of future inflation based on those parameters.

Figure 7.1. Annual inflation rates

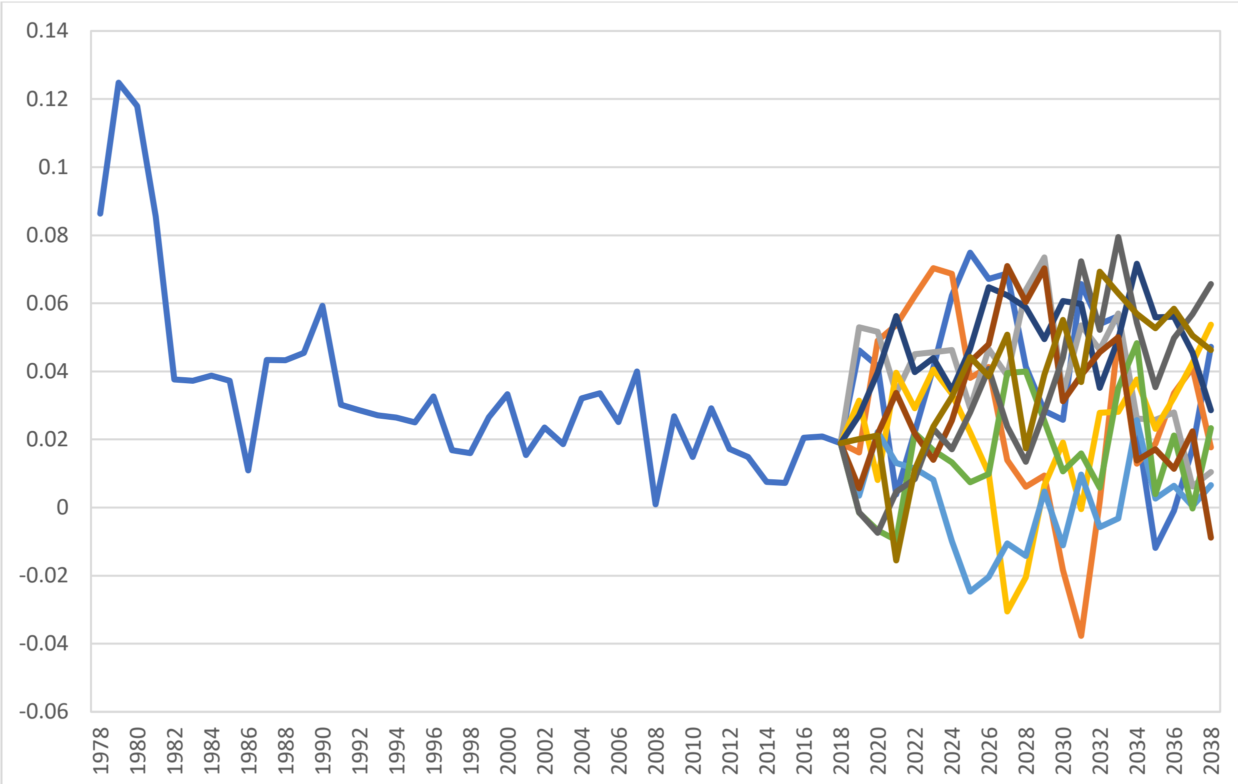

Prepared by Conning, Inc. Sources: ©2020 Bloomberg, L.P. and GEMS® Economic Scenario Generator scenario.

This is conceptually a simple exhibit, yet it provides one with several visual impressions—with some questions. Is the potential future volatility of inflation too high (especially compared with the most recent years of history) when modeled with these parameters? Is the incidence and severity of negative inflation rates too high for comfort? Looked at together, the future projections do initially give the impression of too much volatility. On the other hand, many of the individual simulated inflation rate paths look perfectly reasonable as potential extensions of recent history. We are left with *at least* two questions that we would need to ponder:

- To what degree are we willing to believe that we are now in a different "inflation world" or "regime," and that a low-rate, low-volatility regime will persist for the foreseeable future?

- Although recent global events have suggested that negative interest rates and inflation are real possibilities, just how comfortable are we in permitting our model to simulate such values? And relatively how often?

To continue this visualization process, suppose we choose to take another look at this type of exhibit, but this time with the 1998–2018 data period parameterization from Section 7.2, in the hopes that the simulations may look more representative of the more recent history of inflation rates (Figure 7.2).

Figure 7.2. Annual inflation rates

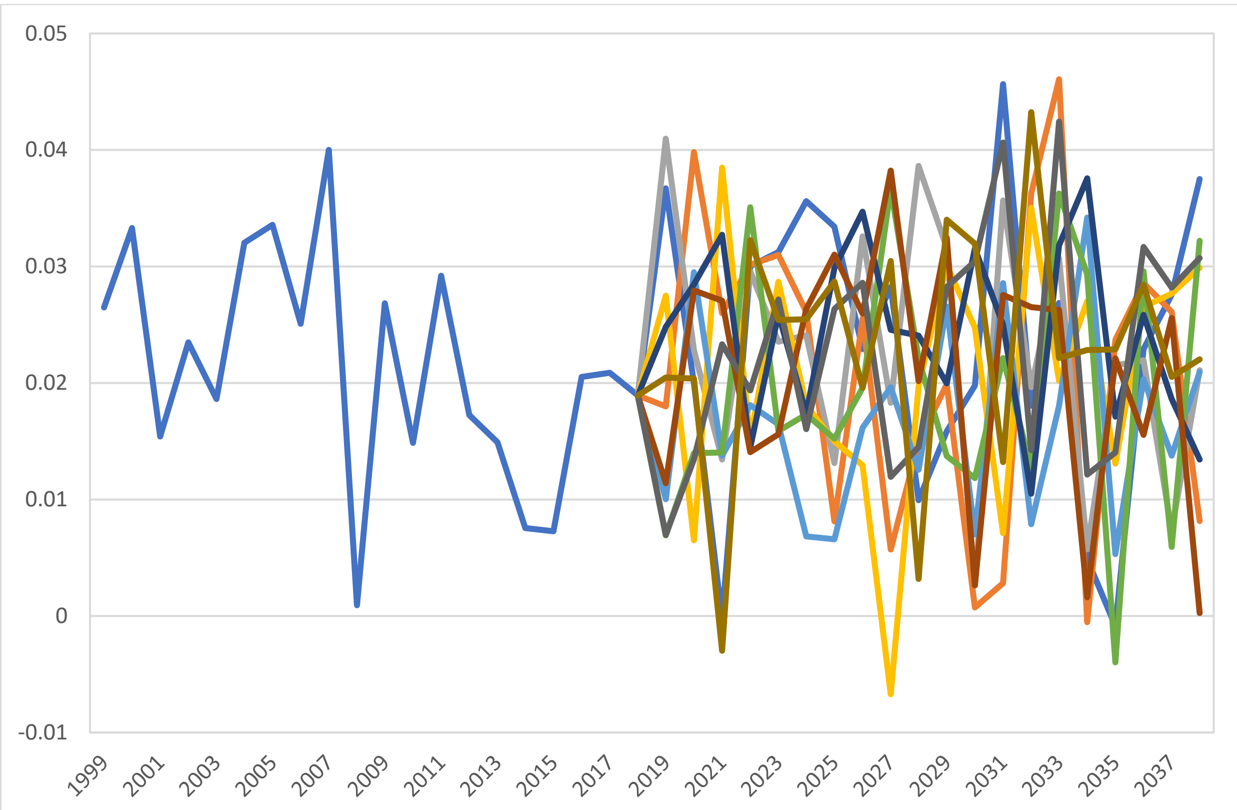

Prepared by Conning, Inc. Sources: ©2020 Bloomberg, L.P. and GEMS® Economic Scenario Generator scenario.

After recognizing that the scale is different in this second exhibit, this does seem to represent the last 20 years of inflation history a bit better than our first exhibit, and certainly the frequency and severity of negative inflation rates have been reduced. However, to accomplish this, we used fewer years of data—which may not be representative of the *long-term* dynamics of inflation.

As always with such analyses, trade-offs are involved between stability and responsiveness. Visual examinations like these can be effective supplements to statistical goodness-of-fit measures and other quantitative assessments.

7.3.2 An Ex Post Observation of Large Stock Returns

Given the gap between when the original ADG model was designed and created and the current data, we have an opportunity to look at the actual emergence of large-stock return data over the last 10 to 15 years to provide an ex post indication of whether that data emergence might be consistent with the original equity model calibration. Has the original model shown robustness with respect to that emerging data?

Another way to ask this question is to examine whether the last 10 to 15 years of additional data have significantly changed the historical distribution of annual stock returns. The following exhibit compares the histogram of large-stock (S&P 500) price appreciation over the period 1872–2005 with that for the period 1872–2018.

Figure 7.3. Histogram of large-stock returns

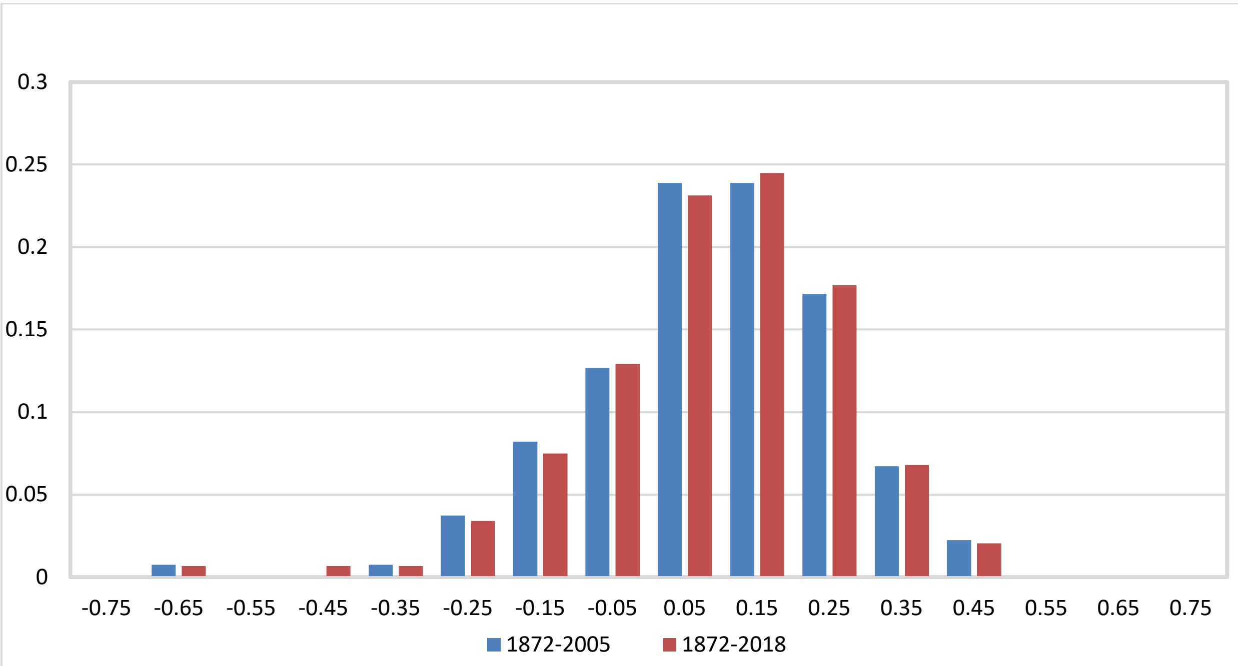

Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.

While quantitative metrics would be important to confirm this, it appears that the last 10 to 15 years of large-stock data have not significantly changed the overall distribution of returns. Visual evidence suggests that, if 15 years ago the ADG model's large-stock return module was producing what were considered to be simulations that represent history reasonably well, that conclusion might still be valid. This is just one of many types of considerations that would help determine the appropriateness of the equity model.

7.3.3 General Comments on Verification and Validation

While the need to verify and validate any model has always been important, if anything it is becoming even more critical. In particular, insurance regulators and rating agencies are emphasizing more and more the need to more closely inspect, assess, and sometimes even approve insurer models. Model assumptions must be vetted, implementation verified, and output assessed for accuracy and reasonableness.

With respect to verification, the goal is to assess consistency between the conceptual structure and the actual model. One can perform such an assessment through several standard procedures that include, among others,

- peer review by third parties;
- clear and sufficient model documentation;
- the creation of graphical representations or flow charts of both the conceptual structure and the actual model, for comparison; and
- varying inputs and checking for the reasonableness of the resulting output and results.

Regarding validation, the goal is to assess the ESG's ability to adequately reflect and represent real-world values, patterns, and behaviors—within the context of the intended purpose and use of the model—and thus provide credibility and confidence for the use of the model. Validation procedures can be either quantitative or qualitative. Perhaps the most common general approach to model validation is the use of a segmented data set, with separate subsets being used for training and for testing/validation, respectively. While some rules of thumb have been suggested as to how much of a data set should be set aside for validation, there is still a fair amount of art involved. A more "informal" type of test might be something comparable to a "Turing test" of model validation. The original Turing test, for machine or artificial intelligence, involves a tester speaking through an anonymity mechanism to both a human and a machine. If the tester cannot tell the difference, the machine can be labeled intelligent. An analogous test for model validation purposes might involve several pieces of data (say, historical graphs of economic variable values, or perhaps descriptive statistics), some real world and some hypothetical and produced by a model. If judges examine these and cannot credibly identify which are which, that might be considered evidence for model validation.

As mentioned, much of the process and the goals of verification and validation depend upon the purpose to which the ESG is applied. As Cairns (2004) says, the accuracy of an aspect of a model (e.g., yield curves) is more important when associated with derivatives pricing and trading than with more general economic/financial analysis. Cairns also recognizes the trade-offs that are inevitable in the assessment of most models: "With the majority of parametric models the aim is to achieve a parsimonious model of the yield curve. Thus, we aim to capture as much as possible of the structure of market interest rates with as few parameters as possible. Clearly, these are conflicting objectives" (pp. 228). One must seek a balance between parsimony and goodness of fit.[28]

7.4 Sample Model Output

In this section, we offer a very brief comment about the ADG model and its output.

The ADG model allows basically any cell in the spreadsheet to be added to the list of outputs and then incorporated into exhibits or charts. Thus, there is literally no projected variable value in the model that cannot be made available as data or in the form of an output exhibit for the user. In

[28] Cairns also considers the use of splines where one wishes to give greater weight to goodness of fit.

fact, as Ahlgrim, D'Arcy, and Gorvett state in their user manual, part of the pre-run preparation when using the model is to define the output cells of interest.

A large number of output options exist—in part, those options are a function of the chart options available in Excel and @RISK. And, of course, simulated data can be downloaded from the ADG model and then used in any other software package for further analysis and exhibit-making.

Figure 7.4 shows an example, taken from Ahlgrim, D'Arcy, and Gorvett (2008), of a type of exhibit the user can automatically create from the ADG model.

Figure 7.4. Ten-year Treasury

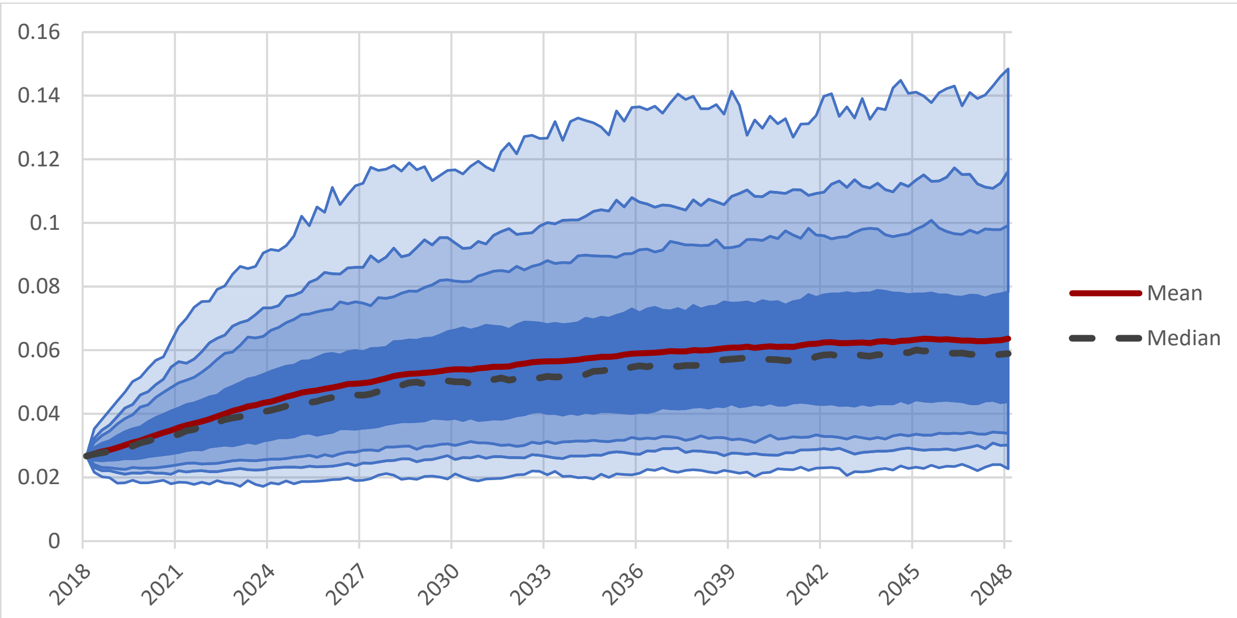

Percentiles plotted: 1%, 5%, 10%, 25%, 50%, 75%, 90%, 95%, and 99%. Prepared by Conning, Inc. Source: GEMS® Economic Scenario Generator scenario.

This "funnel of doubt" chart shows (for projected 10-year nominal interest rates) the growth of confidence intervals associated with simulated values as its projections are made further out into the future. This type of chart can be useful in depicting how the uncertainty of projected results increases with time.

7.5 Some Potential Sources for Economic and Financial Data

Sources of economic and financial data—even of free and publicly available data—are plentiful. Here are a few such sources:[29]

- **Federal Reserve Bank of St. Louis Economic Data (FRED) (https://fred.stlouisfed.org/):** The statement near the top of the FRED homepage (as of this writing)—"Download,

[29] For additional data sources, see Chapter 4.

graph, and track 765,000 U.S. and international time series from 96 sources"—gives an indication of the extent of this economic data repository. FRED is a good general-purpose resource, as it includes data from many sources (e.g., thousands of series from the Bureau of Labor Statistics are included). FRED provides several filters for sorting through this database, not just by type of series but also by geographies, data-reporting frequencies, and specific sources of data. Among the numerous ESG-relevant data series available for downloading are these:

  - *Interest rates*: more than 1,500 series, from three-month T-bills to 30-year T-bonds, and many in between; options include constant-maturity versus secondary-market rates, and daily versus weekly versus monthly values
  - *Inflation*: thousands of series, including the various Bureau of Labor Statistics series that vary by goods, locations, whether or not seasonally adjusted, and data frequencies
  - *Stock prices*: S&P 500, various Wilshire indices, and numerous measures of volatility by geographic location
  - *Measures of economic activity*: including GNP, GDP, and many others
  - *Real estate*: commercial prices, loans, leases
  - *Dividend yields on the NYSE*
  - *Currency and exchange rates*
  - *Commodity prices*

- **United States Department of Labor, Bureau of Labor Statistics (https://www.bls.gov/data/):** Appropriate to the focus of the Department of Labor, this site makes available those time series dealing with key consumer- and employer-related issues. In particular, data of relevance to ESGs found on this site include the following:
  - *Inflation and price levels*: consumer, producer, and international (import/export) price indices
  - *Employment*: including employment numbers, hours, and earnings, as well as other labor force statistics and projections
  - *Unemployment*: national and local unemployment rates and statistics
  - *Pay and benefits*: wages, earnings, and the costs of benefits
  - *Productivity*: by sector and industry
  - *Workplace injuries*: including numbers of injuries and days away from work

- **Yahoo Finance (https://finance.yahoo.com/):** Historical data are downloadable for individual stocks and stock indices.

7.6 Summary

In this chapter, we discuss modeling issues associated with three key ESG variables: inflation, interest rates, and equity returns.

We use the ESG developed by Ahlgrim, D'Arcy, and Gorvett (ADG) as a starting point for our discussion of each variable, but we go beyond the ADG model in this chapter and

- address one or two illustrative issues associated with the parameterization and calibration of a model for those variables; and
- identify some alternative modeling structures for those variables.

The ADG model serves as an illustration of how one research team approached the development of an ESG. The model stems from a joint CAS/SOA RFP and has, from the beginning, been intended for an actuarial/insurance audience. It is publicly available, as is its documentation, and is relatively simple compared with some proprietary ESGs. It is built in Microsoft Excel and is run with the simulation software @RISK, an Excel add-on.

We use the ADG model, as well as its structure and development, to illustrate some of the issues and decisions made with respect to constructing and using an ESG. Since the ADG model is built in Excel, the logic and algorithmic progression of the model is transparent and can be followed by paging through the various sheets of the workbook.

With a publicly available resource such as the ADG model, it is incumbent upon a user to determine whether the parameterization and calibration underlying the model remains adequate. To illustrate, we look at inflation as an example of re-parameterization.

The ADG model uses a continuous-time version of the stochastic mean-reverting process to simulate the dynamics of inflation. Not surprisingly, the parameter estimates for the inflation process vary greatly depending upon the data period used in the regression. If an ESG is supporting short-run operational decision-making, a shorter and more recent data period may better reflect the dynamics of economic and financial variables for that purpose. If long-run strategic planning is the ESG's primary focus, then a longer data period including several economic or business cycles may produce a better range of variable values.

The ADG model employs a double-mean-reverting two-factor Vasicek model for the term structure of real interest rates. The two factors represent the short-term real interest rate r and the long-run mean l. Nominal rates were then determined by combining projections of the real interest rates with those for inflation.

The third key variable, equity returns, is modeled in ADG as a regime-switching process. This regime-switching model, based on the approach of Hardy (2001), involves two regimes, high volatility and low volatility, where the likelihood of staying in the current regime or switching to the other regime is governed by transition probabilities, as in a Markov chain. In ADG, the regime-switching equity returns model is used for simulating *excess* equity returns above a risk-free nominal interest rate. Full equity return rates are then simulated by combining the excess return with the simulated risk-free nominal interest rate. Stochastic volatility models with jumps represent a good alternative to regime-switching models.

The validation process of an ESG is described and discussed in Chapters 4 and 5. In this chapter we provide some brief illustrative examples of possible retrospective validation efforts and some general guidance on the verification and validation processes. We examine simulations of inflation variables and look at ex post indications of emerging stock return data to assess the original calibration of that module. We show the value of visualization and the importance of a holistic view encompassing all modules.

The data used as the basis for the inflation regressions were obtained online from the U.S. Bureau of Labor Statistics. This can be a good source for other economic data such as employment and unemployment, pay and benefits, etc. Other example sources of data include the Federal Reserve Bank of St. Louis Economic Data (FRED) (data on interest rates, inflation, measures of economic activity, real estate, currency, etc.) and Yahoo Finance (downloadable data for individual stock and stock indices).

References

Ahlgrim, K. C., S. P. D'Arcy, and R. W. Gorvett, "Report on Modeling of Economic Series Coordinated with Interest Rate Scenarios," Arlington, VA: Casualty Actuarial Society, 2004, https://www.casact.org/research/econ/.

Ahlgrim, K. C., S. P. D'Arcy, and R. W. Gorvett, "Modeling Financial Scenarios: A Framework for the Actuarial Profession," *Proceedings of the Casualty Actuarial Society* 92, 2005, pp. 177–238.

Ahlgrim, K. C., S. P. D'Arcy, and R. W. Gorvett, "A Comparison of Actuarial Financial Scenario Generators," *Variance* 2:1, 2008, pp. 111–134.

Cairns, A., 2004, *Interest Rate Models: An Introduction*, Ch. 12, Princeton, NJ: Princeton University Press, 2004.

Christensen, I., F. Dion, and C. Reid, "Real Return Bonds, Inflation Expectations, and the Break-Even Inflation Rate," Working Paper 2004-43, Ottawa: Bank of Canada, 2004.

CIA (Canadian Institute of Actuaries), "Calibration of Equity Returns and Volatility for Stochastic Models," Research Paper, Document 217055, Ottawa: CIA, 2017.

Fama, E. F., "The Behavior of Stock Market Prices," *Journal of Business* 38:1, 1965, pp. 34–105.

Gürkaynak, R. S., B. Sack, and J. H. Wright, "The TIPS Yield Curve and Inflation Compensation," Finance and Economics Discussion Series Working Paper 2008–05, Washington, DC: Federal Reserve Board, 2008.

Hardy, M. R., "A Regime-Switching Model of Long-Term Stock Returns," *North American Actuarial Journal* 5:2, 2001, pp. 41–53.

Hardy, M. R., R. K. Freeland, and M.C. Till, "Validation of Long-Term Equity Return Models for Equity-Linked Guarantees," *North American Actuarial Journal* 10:4, 2006, pp. 28–47.

Heath, D., R. Jarrow, and A. Morton, "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation," *Econometrica* 60:1, 1992, pp. 77–105.

Hibbert, J., P. Mowbray, and C. Turnbull, "A Stochastic Asset Model & Calibration for Long-Term Financial Planning Purposes," Technical Report, Barrie & Hibbert Limited, 2001.

Hull, J., and A. White, "Pricing Interest-Rate Derivative Securities," *Review of Financial Studies* 3:4, 1990, pp. 573–592.

Vasicek, O., "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics* 5, 1977, pp. 177–188.

Appendix 7.A

Flow chart of ADG model economic variable generation

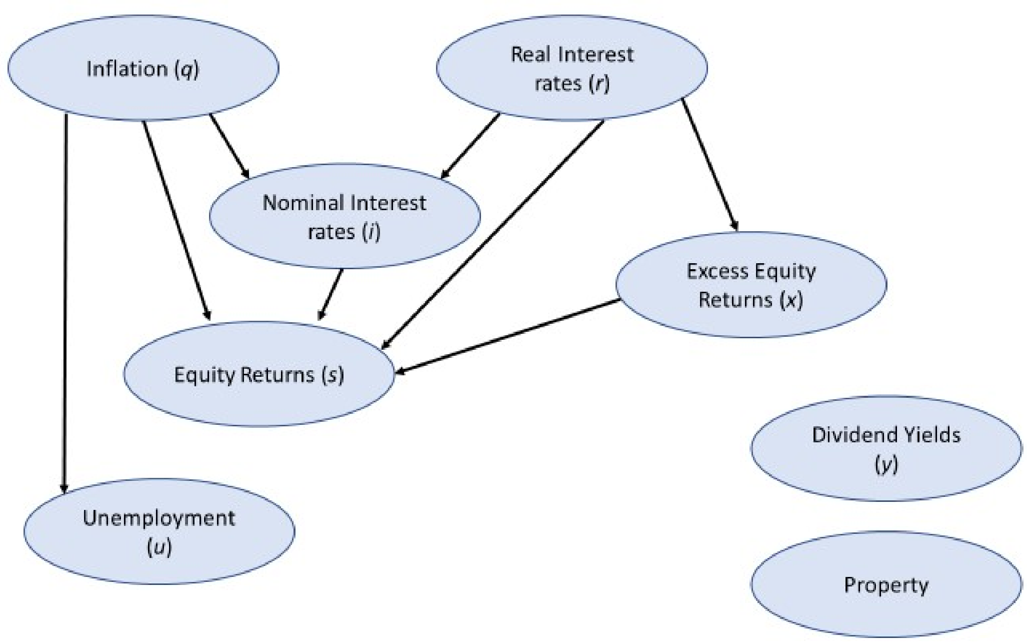

Appendix 7.B

Summary of ADG model stochastic process descriptions and regression equations

Inflation	q	$dq_t = \kappa_q(\mu_q - q_t)dt + \sigma_q dB_q$ Mean-reverting process.	$q_{t+1} = \alpha + \beta q_t + \varepsilon'_{qt}$ Regression on annual data.
Real interest rates	r	$dr_t = \kappa_r(l_t - r_t)dt + \sigma_r dB_r$ $dl_t = \kappa_l(\mu_r - l_t)dt + \sigma_l dB_l$ Mean-reverting to a long-run mean-reverting mean.	$r_{t+1} = \alpha_1 l_t + \alpha_2 r_t + \varepsilon'_{rt}$ $l_{t+1} = \beta_1 + \beta_2 l_t + \varepsilon'_{lt}$ Plus two-stage least squares est. regressions on monthly data.
Nominal interest rates	i	$P^i(t,T) = P^r(t,T) \times P^q(t,T)$ Bond prices from underlying inflation and real rate processes.	
Equity returns	s	$s_t = q_t + r_t + x_t$ x: excess equity return attributable to capital appreciation.	Low- and high-volatility regimes. Small and large stocks.
Equity dividend yields	y	$d(lny_t) = \kappa_y(\mu_t - lny_t)dt + \sigma_y dB_{yt}$	Estimation is analogous to the process used for inflation.
Real estate (property)	re	$d(re)_t = \kappa_{re}(\mu_{re} - (re)_t)dt + \sigma_{re} dB_{re}$	Based on data from National Council of Real Estate Investment Fiduciaries.
Unemployment	u	$du_t = \kappa_u(\mu_u - u_t)dt + \alpha_u dq_t + \sigma_u dB_{ut}$	$u_{t+1} = \beta_1 + \beta_2 u_t + \beta_3(q_{t+1} - q_t) + \sigma_u \varepsilon'_{ut}$

Chapter 8: Considerations Related to the Projection Time Frame (Simulation Horizon)

In this chapter, we cover some of the considerations that arise when using an ESG for projections over a specific time period ("simulation horizon"). Such simulation horizons can be quite short, e.g., one year, which is typical when modeling economic capital. However, much longer simulation horizons are not uncommon, such as, for example, 10, 20, or 30 years for asset–liability management (ALM) or strategic asset allocation (SAA) purposes. Since the focus of the application will define how the ESG is used, it is useful to know what aspects of the ESG calibration one needs to consider in these different cases. In the case where one receives ESG scenarios or calibrations from an external provider one should be clear about the intended purpose of the scenarios or calibrations and the corresponding projection horizon as well.

In this chapter, we concentrate on the following:

- Section 8.1: General considerations in selecting an ESG for a specific simulation horizon
- Section 8.2: Mean reversion, return, and correlation properties over different projection horizons
- Section 8.3: Important considerations in generating ESG scenarios over a long-term horizon
- Section 8.4: Important considerations in generating ESG scenarios over a short-term horizon
- Section 8.5: Challenges in generating coherent scenarios for multiple simulation horizons

We start with some general aspects of ESGs when they are used for a specific simulation horizon. In Section 8.2 we take a more detailed look at some of those aspects. In Sections 8.3 and 8.4 we look at the issues to take into account when simulating for shorter or longer time horizons. In the last section, we consider challenges arising when one needs to generate coherent scenarios for multiple simulation horizons.

8.1 General Considerations in Selecting an ESG for a Specific Simulation Horizon

Depending on the ESG application, the calibration targets will be different. As mentioned, whereas risk-based applications will focus more on the tail aspects of the distribution, applications such as ALM or SAA will put more focus on the central properties of the distribution. However, "focusing" in this context does not mean that other distributional properties can be completely forgotten. Rather, it is a way of ordering what aspects come first in the validation of ESG results and what aspects are of secondary importance.

Choosing the specific simulation horizon for an ESG is tightly connected with the overall use to which the ESG scenarios will be applied. Quite short horizons (e.g., one year) can be seen for economic capital simulations that analyze the risks of an insurance company. Such simulations are more focused on the tail aspects of the distributions, like percentiles and value at risk. Longer simulation horizons are more common in applications such as ALM or SAA. The results of these applications are centered around the central properties of the distributions, like mean and standard deviation.

However, the distinction between short and long simulation horizons and where in the distribution they are focusing can be somewhat blurry. For example, one-year ALM projections or three-year risk projections can also be encountered in practice. In the end, it is the application of the ESG that decides what calibration targets are appropriate.

For example, let's consider a short-term risk application of the ESG. Here one could set quite strong targets for the percentiles of 10-year interest rates after one year, i.e., enforcing low tolerance levels for deviations from target tail percentiles, and after that also set targets for the mean but with a higher tolerance level.

One should also keep an eye on the current market conditions and be clear that those are also likely to have an effect on the calibration targets. So, for example, if the shape of the starting term structure curve is unusual (e.g., negative 10 year–1 year spread) one might expect that this shape would not stay as it is throughout a long simulation but rather (slowly or quickly) evolve to a more regular shape. Since we are talking about stochastic scenarios here, the user's expectation may be that at some point in the projection period the *majority* of the yield curves generated by the ESG will have regular shapes (i.e., some tail event curves may still have a negative 10 year–1 year spread).

Note that for short-term, risk-focused simulations the targets for the interest rate percentiles should ideally depend on the overall interest rate level. So, for example, the probability of a 10-year rate going down 150 basis points should be higher for the cases when the interest rate is high, e.g., circa 4%, and lower for the cases when the interest rate is low, e.g., circa 1%.

Implied market volatilities can also play a role in the ESG calibration. Let us consider a long-term ALM simulation. Given a situation where the current implied equity volatilities are rather high, one would prefer that the realized equity volatilities be higher at the start of the simulation and then revert over time to the long-term equity volatility levels. This is more logical than the situation where the realized equity volatilities do not react to the starting market conditions.

8.2 Mean Reversion, Return, and Correlation Properties over Different Projection Horizons

In this section we carry out a comparison of some of the ESG outputs for a calibration as of December 31, 2018 (base scenario) and a stressed calibration that starts with a 400–basis point higher yield curve (stress scenario). We change only the ESG parameters responsible for the starting curves, leaving the long-term parameters unchanged.[30] Therefore, we will see the effects

[30] The interest rate model used in this section is a three-factor Cox–Ingersoll–Ross model, calibrated to 60 years of U.S. historical data.

of the different starting conditions as well as the effects of the mean reversion in the graphs that follow.

The graph in Figure 8.1 shows the development of correlations between returns of two-year and 10-year bonds as the simulation progresses.

Figure 8.1. Correlations between two-year and 10-year bond returns

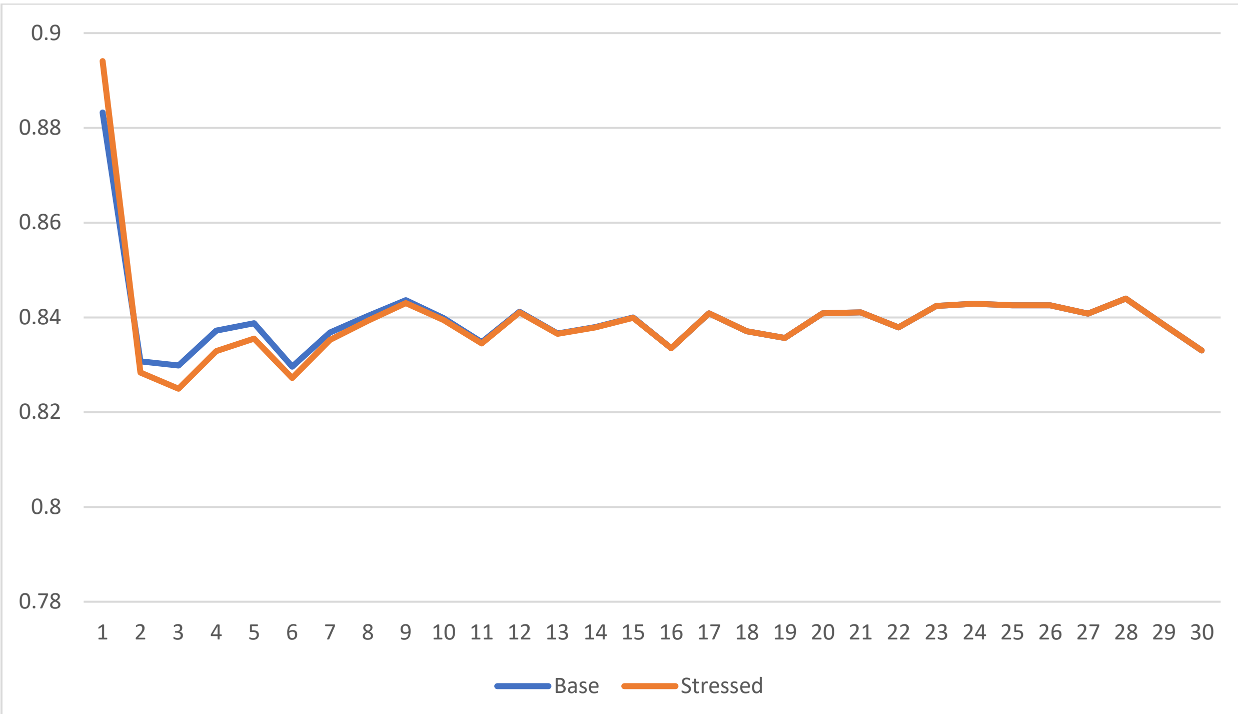

Prepared by Conning, Inc. Source: GEMS® Economic Scenario Generator scenario.

As one can see, the correlations are not constant throughout the simulation horizon but instead evolve over time. Also, after some years the correlations converge to the same values regardless of the starting yield curve. This is because both ESG calibrations use the same long-term parameters.

We see a similar—though somewhat different—development on the next graph (Figure 8.2), where we analyze the returns of a 10-year bond.

Figure 8.2. Ten-year bond returns

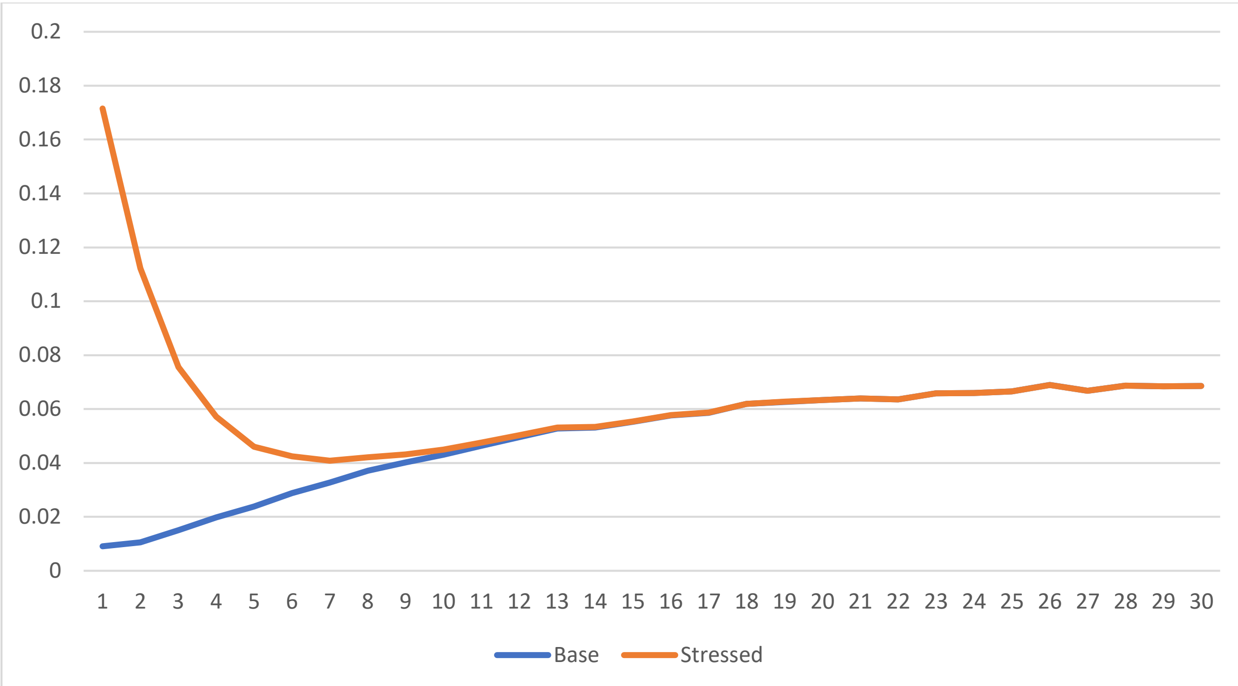

Prepared by Conning, Inc. Source: GEMS® Economic Scenario Generator scenario.

The much higher bond returns in the stressed case are explained by two factors: (a) generally the yields are higher in the stressed case, and (b) the yields are decreasing over time in the stressed case, thus increasing the bond returns in general. Here, again, we see the differences introduced by the starting conditions, but such differences disappear over time as the simulation progresses. Again, both ESG calibrations converge to the same results because we use the same long-term parameters.

8.3 Important Considerations in Generating ESG Scenarios over a Long-Term Horizon

When running an ESG over a long-term horizon the stability of the results is quite important to the user. Ideally, long-term decisions should not be significantly affected by short-term market changes. The best way to achieve this stability is to calibrate the ESG to stable long-term targets. Those targets will typically be the long-term levels of interest rates (mean and standard deviations) and equity returns (mean and standard deviations).[31] This is because without long-term targets in the ESG, one should not expect any result stability over a long-term horizon.

[31] In an SAA context, one could also ask the following question: "Under what economic conditions would the current investment strategy be optimal?" However, it must be noted that the search for an answer in this case can be quite long, and such economic conditions may not be found at all under certain circumstances.

One can establish the long-term targets in the ESG either by using an expert view from within one's company or by using the historical data. However, one needs a certain level of caution when working with historical data. For example, if a chosen data series is too short, it could contain some trends that one would rather not have in a simulation or the distribution of the interest rates could be too narrow in the historical data. Choosing a longer historical data series doesn't guarantee that all possible issues will be resolved because the longer time series may, for example, contain extremely high interest rates that may be less likely to happen seen from today than suggested by the historical data alone. Therefore, blindly following the historical time series may not be the best course of action when using historical data for setting the long-term simulation targets. We will encounter a similar situation when we cover the setting of short-term targets.

The speed of convergence to the long-term targets can also be important, since this will determine how much influence the starting conditions will have on the overall simulation results. Depending on the convergence speed, the results of the simulation can be quite different even for the same starting point and the same long-term targets. This is illustrated in the graph in Figure 8.3, which shows a potential mean path for future interest rates.

Figure 8.3. Example of different conversion speeds

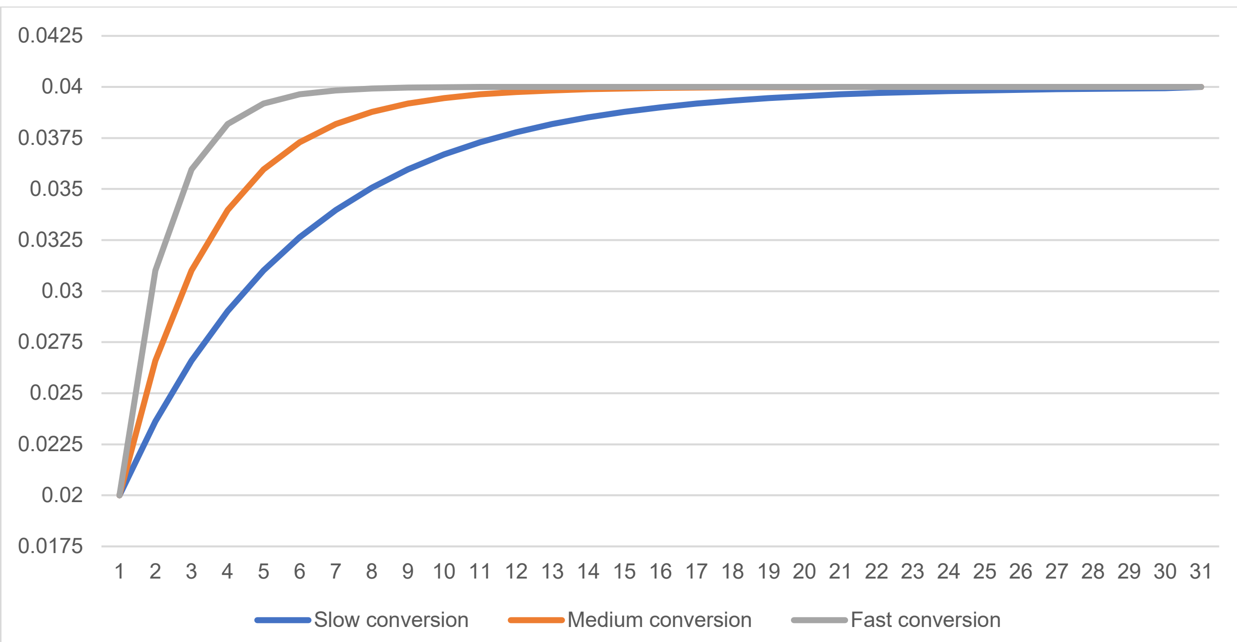

Prepared by Conning, Inc. Source: GEMS® Economic Scenario Generator scenario.

Another important aspect to consider when running ESGs for long-term horizons is whether the downstream system can handle unexpected results produced by the ESG. Examples of such unexpected results include negative yields or very high yields, above 100%, which some of the interest rate models may produce when run for longer time periods.

Short-term considerations should be treated quite carefully when calibrating ESGs to long-term horizons. On the one hand, the starting points of the interest rate curves are determined by market data. On the other hand, the overall results of applying the ESG should ideally be stable. However, very rapid average interest rate movements to the long-term targets, in the case of fast conversion, can cause quite large and unexpected bond returns in the first years of the simulation, where especially the results in the first year can be critical. So, the user should be prepared to answer questions around the expected bond returns not only for the long-term targets but also for the starting years of the simulation. A user should also make sure that the returns generated by the ESG are in line with the user's acceptance criteria. The latter can be done when validating the ESG results.

8.4 Important Considerations in Generating ESG Scenarios over a Short-Term Horizon

When running ESG scenarios over a short-term horizon a common application is risk-focused analysis. Therefore, the tails of the distributions will be more important when calibrating the ESG for such an application. However, it is quite common to encounter in practice percentile targets that are based on quite rare events, like the 0.5th percentile that represents a 1-in-200-year event. Several challenges arise when establishing calibration targets in this case. On the one hand, the historical data may not be long enough to establish a statistically sound target and expert judgment may be required. On the other hand, some of the historical data may refer to time periods with a very different market situation, which, therefore, can be of no or very limited use for establishing the calibration targets.

For example, if we consider setting the 0.5th percentile target for the movement of 10-year interest rates we could question the applicability of market data from the first half of the 1980s, where such rates were in double digits and one year movements of more than 300 basis points occurred. Of course, one could argue that such historical data could be used to establish targets for relative changes in the yields. However, even that approach could be questionable. For example, such relative targets, when calculated for yields of 2% or 3%, which are quite common nowadays, will actually differ by 50% when converted to absolute yield changes. Such large differences in the percentile targets for fairly similar initial market conditions appear unlikely and are rather hard to explain.

Having established the targets for a risk-focused calibration, one still needs to monitor changes in market conditions to ensure that the targets remain appropriate. Two types of market movements may make adjustments to the targets necessary: (a) a major market movement that changes the assumption of what a 1-in-200-year event may be (an example of that was seen in 2008, when the jumps in credit spreads set new historical records for spread movements); or (b) a gradual market movement that slowly but surely introduces a new economic situation to the market's starting conditions (for example, the gradual reduction in interest rates over the last couple of decades). The annual yield changes in the latter example do not, individually, represent an event at the extreme end of the interest rate movement distribution, so the first criteria for reassessing the calibration targets will not be triggered. However, over time the combination of incremental movements is significant, and it should be clear that the potential downside risks should be higher in the case where the starting interest rate is 5% as opposed to when it is just 2%.

If the calibration targets are set and they are stable, there is, of course, still the question of the stability of the resulting model parameters. Ideally, the model parameters should stay quite stable if the initial market conditions do not change much. However, since in some interest rate models the interest rate volatility is proportional to the starting level of the interest rates, the parameters can change quite a bit even if the interest rate levels do not change much. In a very simple example,[32] the interest rate volatility is just $\sigma * r_0$, where r_0 is the starting interest rate. Assuming that r_0 is initially 2% but changes to 1.8%, which is a rather small absolute movement in the interest rate, the new sigma should increase by circa 11% to obtain the same level of interest rate volatility. Therefore, one should not necessarily expect very stable model parameters even if there are no changes in the calibration targets, as the changes in model parameters can be a result of (small) changes in the initial market conditions.

It should also be clear that even if the calibration targets are stable differences can still occur in the returns of the underlying assets. For example, if the 10-year zero-coupon yield moves 100 basis points from 1% to 2%, the resulting change in the zero-coupon bond price is circa −9.38%. If, however, the same movement happens from 2% to 3%, the resulting change in the zero-coupon bond price is circa −9.30%. Although that is a minor difference, it highlights the fact that even if the distributions of the interest rate movements are identical between different ESG calibrations, the resulting return distributions can be different due to differences in the starting curves.

Therefore, we have two possible change drivers in an ESG used for short-term risk-focused calibration: (a) change in targets/calibration of the ESG, and (b) change in the starting market conditions. Whereas the first change driver is easy to explain and to observe, the second change driver can be more subtle.

8.5 Challenges in Generating Coherent Scenarios for Multiple Simulation Horizons

An interesting challenge arises when one wants to unite both short-term and long-term criteria in one ESG scenario. In such a case, ideally the ESG should generate percentile-focused results at the start of the simulation and mean- and standard deviation–focused results some time into the simulation.

The difficulty here lies in the fact that the distributions one is focusing on in the short-term can be quite different from the ones of interest in the long-term simulations, so just taking the same distribution and applying it independent of the level of the random variable that is simulated will not lead to the desired outcome. On the other hand, in some stochastic interest rate models the volatility of the interest rate is proportional to the level of the interest rates or to the square root of the level. Here it is possible to obtain more volatility when interest rates are higher. However, in certain cases—especially when the starting point is rather far away from the long-term target—

[32] Here we consider a one-time step model with $r_1 = r_0 + \alpha + \sigma * r_0 * N(0,1)$, where $N(0,1)$ is the standard normal random variable, and α and σ are constant.

this could lead to undesired effects. Let us consider, for example, a very simple case where the starting point lies at 2% and the long-term target is at 4%, and a model where the level of interest rates is the driver for the interest rate volatility. Applying the same model parameters for the starting point and the long term will result in an interest rate volatility that is two times higher for the long term than it is for the starting point. Just a reminder: if the underlying distribution in the change of interest rates is normal, that would actually mean that the movements to the quantiles are now double what they were at the start.[33] Such a sharp increase in the interest rate volatility could be excessive and may not withstand validation against historical data.

For a model where the square root of interest rates is the driver of volatility, consider an example where the starting point is 1% and the long-term level is 4%. This leads to the same effect of doubling the interest rate volatility, and, again, it could result in either too much or too little volatility in the long term, depending on the exact targets that are set for the start of the simulation.

The models that are used in practice are usually a bit more sophisticated, but they can be separated into two broad classes: (a) where the volatility of interest rates is not dependent on the level of interest rates, and (b) where the volatility of interest rates rises with the level of interest rates. For the first class of models, it is impossible to construct an ESG where both short-term and long-term calibration criteria will be met. For the second class of models, there is a risk that the model parameters, when calibrated to the short-term criteria, will either overestimate or underestimate the long-term distribution parameters. However, there is at least a possibility of finding a suitable compromise between short-term and long-term calibration criteria in this class of models.

[33] At the start there is also some movement to the mean reversion level; however, it is typically rather small compared with the increase in the volatility we are discussing here.

Chapter 9: Calibrations for One-Year and Short-Horizon Capital Models

In some circumstances an ESG user needs a model designed and calibrated specifically for the projection of assets and liabilities at short simulation horizons. These applications include regulatory capital calculations, such as those required under the U.S. risk-based capital framework and the European Solvency II Directive's Own Risk and Solvency Assessment, as well as capital models used for general risk management and investment decision-making. Defining and implementing a calibration of this type poses some specific challenges.

For the purposes of this discussion we will assume that "short term" refers to a simulation time horizon of between one year (typical for regulatory reporting) and three years (more typical for general risk and investment decision-making). While this choice may not encapsulate every conceivable use of a model of this type, it captures the horizon that best illustrates some of the practical aspects of an implementation.

This chapter introduces the key aspects one must consider when implementing a short-term calibration and discusses in general terms some of the more common approaches.

9.1 Short-Term Calibrations—Important Considerations

A modeler needs to address several questions prior to the specification and implementation phase of a short-term calibration. Broadly speaking, one must consider three areas at the outset: first, target-setting or the establishment of suitable benchmarks to ensure that the model is fit for purpose; second, calibration or parameterization, the task of finding appropriate model parameters that satisfy our benchmarks; and third, maintenance, which determines what must be done to ensure the benchmarks and models remain appropriate and relevant.

Benchmarking first requires the user to define specifically a set or subset of variables and statistical properties for which targets should be set. This may start with a list of economic variables for which the short-term calibration is relevant; interest rates, equity returns and dividends, other risk assets (e.g., alternatives such as private equity), credit spreads, inflation, and foreign exchange are typical in an exercise of this kind. These economic variables may then need to be broken down into more granular classifications by, for instance, region (e.g., U.S. equity and foreign equity) and, where appropriate, credit rating classes (e.g., investment grade, high yield, or more granular). Finally, the user must define for which statistics targets or benchmarks must be set. This may depend on the usage. For instance, for regulatory capital applications, a high degree of focus, as well as impact on the final capital numbers, is concentrated on tail metrics such as the 0.5th, 1st, 99th, or 99.5th percentile. For all applications, setting targets for the standard deviation as well as the best estimate of a variable are likely important. One must also consider if the best estimate is equivalent to the mean or the median of the projected variable. Correlations may also be important, as may higher-order distributional properties such as the skewness and excess kurtosis, and such importance should be assessed in some way.

The models or ESG one chooses for the short-term calibration may depend on what has been decided in the benchmarking stage. For instance, if a user has decided that tail measures are an important aspect of the benchmarking procedure, then a model and calibration environment that specifically allows for the calibration to tails is important. Questions one must ask at this stage of planning include these: Is the model capable of producing distributions that match the benchmarks? How much deep knowledge of the models is required to calibrate them? Do we need to build the calibration environment ourselves? What validation and acceptance criteria are needed?

With appropriate benchmarks in place, we next need to calibrate the model. We could do this by using the benchmarks directly as specified targets and constraining the model parameters to values that match the targets closely. Alternatively, we could calibrate the model using analytical algorithms run on time series market data (e.g., as in Feldhütter 2016) using the benchmarks for validation of the output. In practice, however, although the latter option may lead to reasonable fits to historical distributions of data it can potentially produce short-term simulation results that may not align well with actual observed market behavior. For instance, calibrating a model using time series data on interest rates may lead to a model with unreasonably high expectations of interest rates in the long term with associated excessive mean reversion, reflecting more the historical mean than likely short-term behavior. For this reason, most short-term calibrations are formed by fitting the model directly to the benchmarks as targets or using analytical estimation (i.e., using time series data) heavily constrained on those benchmarks.

We must also consider what ongoing maintenance of the calibration we might need. Short-term benchmarks are by nature sensitive to current or initial conditions and one can reasonably expect them to react more strongly to new information in the market than longer-term calibration benchmarks. In many cases, the models must be calibrated precisely to the initial market conditions, which are almost always dynamic. The process for updating the benchmarks and the models based on such factors is also important.

In summary, no consensus or blueprint exists for how one should implement a short-term calibration; what is important is that the usage of the ESG and the calibration is well understood and that this influences the process of defining, implementing, and maintaining a calibration.

9.2 Establishing Calibration Benchmarks

Establishing short-term calibration benchmarks is arguably more difficult than calibrations for the longer-time-horizon case. Part of this is because significant variability is observed in the distributions and distributional properties of financial variables over different short time horizons. Take, for example, the distributions of daily equity price returns for the S&P 500 index separately for the years 2001, 2008, and 2013, as shown in Gouriéroux. We observe quite different distributions, with 2013 being a high-return/low-risk year, while 2008 is markedly more volatile, with fatter, longer tails on both the left and right side. And 2001 sits somewhere in the middle.

Figure 9.1 U.S. large-cap daily equity return by year

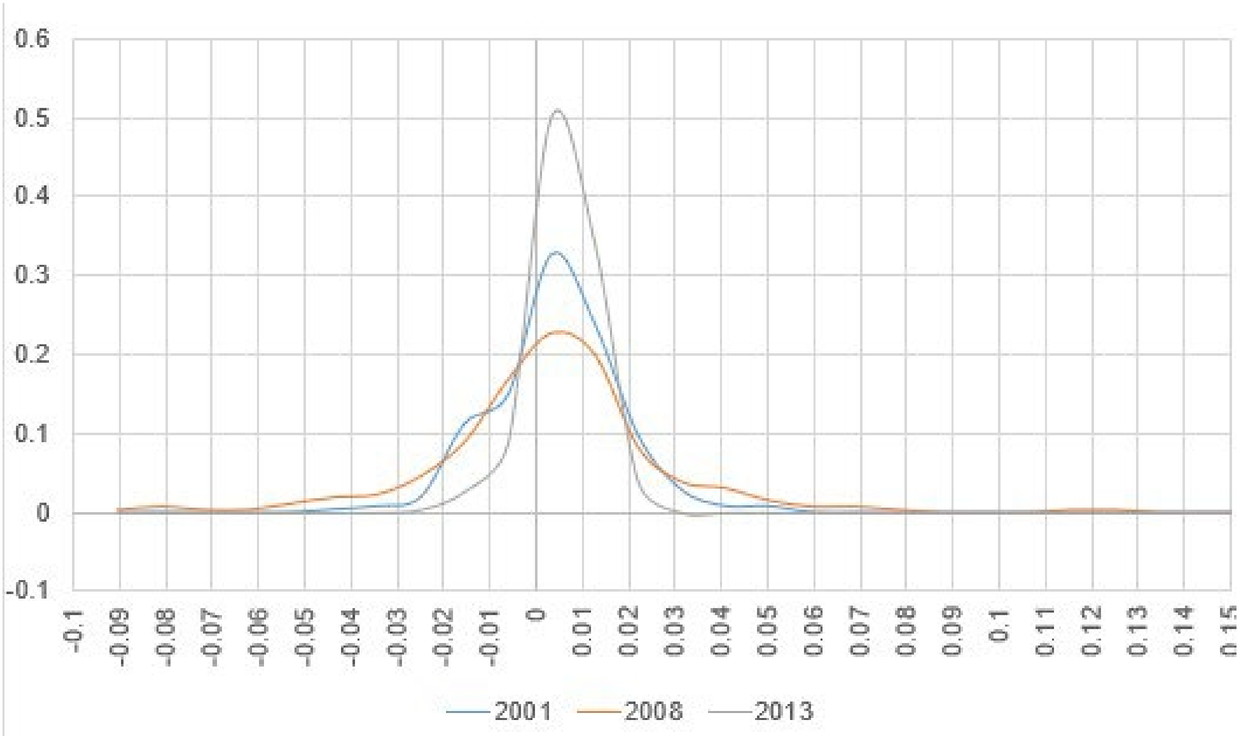

Distribution of daily equity returns for the S&P 500 in 2001, 2008 and 2013. Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.

We can gain some more clarity on the similarities and differences in these return distributions by looking at the statistical properties of the above distributions shown in Table 9.1.

Table 9.1. Statistical properties of U.S. large-cap daily equity returns

Statistic	2001	2008	2013
Mean	-0.05%	-0.16%	0.07%
Std dev	1.36%	2.57%	0.73%
0.50%	-4.11%	-8.89%	-2.30%
1%	-3.28%	-8.14%	-2.30%
Annual return	-13.04%	-37.58%	29.60%

Statistical properties of daily equity returns for the S&P 500 in 2001, 2008, and 2013. Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.

As the distributions imply, the statistical properties of daily returns in 2001, 2008, and 2013 were quite different. The volatility of returns was approximately 3.5 times larger in 2008 than in 2013, and the lower percentiles were also significantly worse. Comparing the price return over the whole year, we see that the return ranges from −37.58% to +29.60%. This all serves to highlight the dynamism of financial markets and that we cannot hope to accurately forecast or reflect the distributions of a variable over a short time horizon. Instead, we must aim for a prudent and robust view to reflect in our models.

Two considerations are important when establishing calibration benchmarks. The first is which risk metric or metrics we should focus on during the target-setting procedure (e.g., volatility or some percentile or percentiles). The second is the question of what methodology or procedure we should use to set the numerical value of the benchmark statistic for a given economic variable at the chosen time horizon.

Most applications will require calibration benchmarks for all or some of the following:

- *Initial conditions* (often called time-zero values), which act as a starting point for all simulations
- A *best estimate*, which most typically corresponds to the mean or median of the simulated distribution
- A *volatility benchmark*, giving a target for the standard deviation of the simulated distribution
- *Relevant extremes* that the simulated distribution should exhibit

We now look at the important aspects of each of these.

9.3 Time-Zero Values

The time-zero values are important because they ensure that all simulations start from a realistic point and that asset and liability valuations in the model reflect the true book value of the firm. These time-zero values almost always come from a third-party data vendor and correspond to the market conditions at the valuation date. This may not be true for alternative and illiquid asset classes, however, where a value may need to be extrapolated or obtained from the issuing party. There are also cases where a yield curve or other market variable must be extrapolated beyond the last liquid point. In that case, the user must define a suitable extrapolation technique or, in the case of some regulatory regimes, use the designated method or published curves from the regulatory body.

It is also worth mentioning that the availability of time-zero data is often one of the biggest constraints on delivery timelines for a calibration of this sort. Some types of data are often made available with several days of lag to a month end. It is important to understand this at the outset and before committing to a particular delivery timetable.

9.4 Estimation of the Best-Estimate Mean Benchmark Value

For the best-estimate mean (or median), a user has several possible ways at his or her disposal to set the benchmark for each economic variable. For instance, any one of the following might lead to a reasonable or justifiable target for a given financial or economic variable:

- **Historic data analysis:** One typically uses short (e.g., 1, 2, or 5 years of data) histories of data for short-term targets.
- **Economic assumptions:** One uses central bank information or forecasts as a basis for setting targets.
- **Zero drift: One** assumes that the market in the short term will stay where it is today.
- **Forward rates:** One uses market-implied forward rates as the basis of short-term forecasts.
- **Econometric forecasting models:** One uses econometric or other modeling techniques (e.g., machine learning) to forecast a variable and use that as the basis of a target for the mean.
- **Analyst or investment committee view:** One sets the model mean to the company investment viewpoint or that of a committee of analysts or other stakeholders.

Each method has pros and cons, and the one a user chooses should depend in part on the economic variable under consideration.

Historical data analysis is arguably very useful, particularly where we may have periods of stable and stationary data that point to the most probable outcome in the future. However, analysis of historical data should be used with some caution. Where data exhibits strong trending, as for instance in some interest rate environments in the decade post-2008, one must carefully consider whether the measured mean and standard deviation have statistical or intuitive meaning as it pertains to the benchmarking exercise. Also, where the economic environment has changed structurally for some reason, we must consider which history is most relevant or how much weight to apply to particular periods of history. This is the case, for instance, when considering inflation data where levels and volatility of inflation rates in many economies have been observed to decrease following the adoption of central bank targeting in the 1990s. Either way, the user must almost always employ expert judgment in using historical data to create a benchmark for the mean of a short-term calibration.

Many central banks and intergovernmental organizations, such as the Organisation for Economic Co-operation and Development and the World Bank, publish specific guidance on future expectations regarding certain financial variables. For some variables, such as short-term interest rates and inflation, one can employ that guidance as the basis of a target mean. This has several advantages—in particular, the guidance may be published by the body directly responsible for setting the future value of the variable of interest or that at least has a strong influence over it. Also, from a process point of view, it is simpler to devolve the benchmarking exercise to a credible external body than to manage the process internally. Disadvantages to this approach might be that one often has little in-depth insight into how the target was arrived at and that the benchmark may change significantly from month to month.

The "zero-drift" assumption—i.e., assuming that a given variable remains on average at its time-zero value for the entirety of the simulation—may be appropriate for some variables and some applications. For instance, in a short-term capital model used for strategic investment decision-making, it may be useful and appropriate to assume that foreign exchange rates do not drift or mean-revert to some value in the short term. This is because the analyst may want to consider what the optimal strategy is in the absence of any bias caused by positive or negative returns attributable to foreign exchange effects. For other variables such as equity, however, this approach is clearly inappropriate. No one would argue that, because the S&P 500 price return for 2008 was −37.58%, the benchmark mean for 2009 should also take that value.

Using the forward rate or the forward price for a traded market instrument like the use of central bank targets has the advantage of devolving the benchmarking exercise away from the analyst. In this respect it is, from a process point of view, simple to implement and requires little additional analysis. The approach may also be relatively easy to justify, in the sense that it can always be argued that one is adhering to the market view. One drawback, however, is that the forward rate is quite volatile, as can be seen in Figure 9.2, which shows the historical 12-month forward rate for the U.S. dollar–to–British pound foreign exchange rate. Also, we observe in Figure 9.2 that the forecasting power of this forward rate is quite poor, tending to reflect the initial conditions of the foreign exchange spot rate very closely rather than matching future spot rates.

Figure 9.2. USD/GBP forward rate and USD/GBP spot rate lagged 12 months

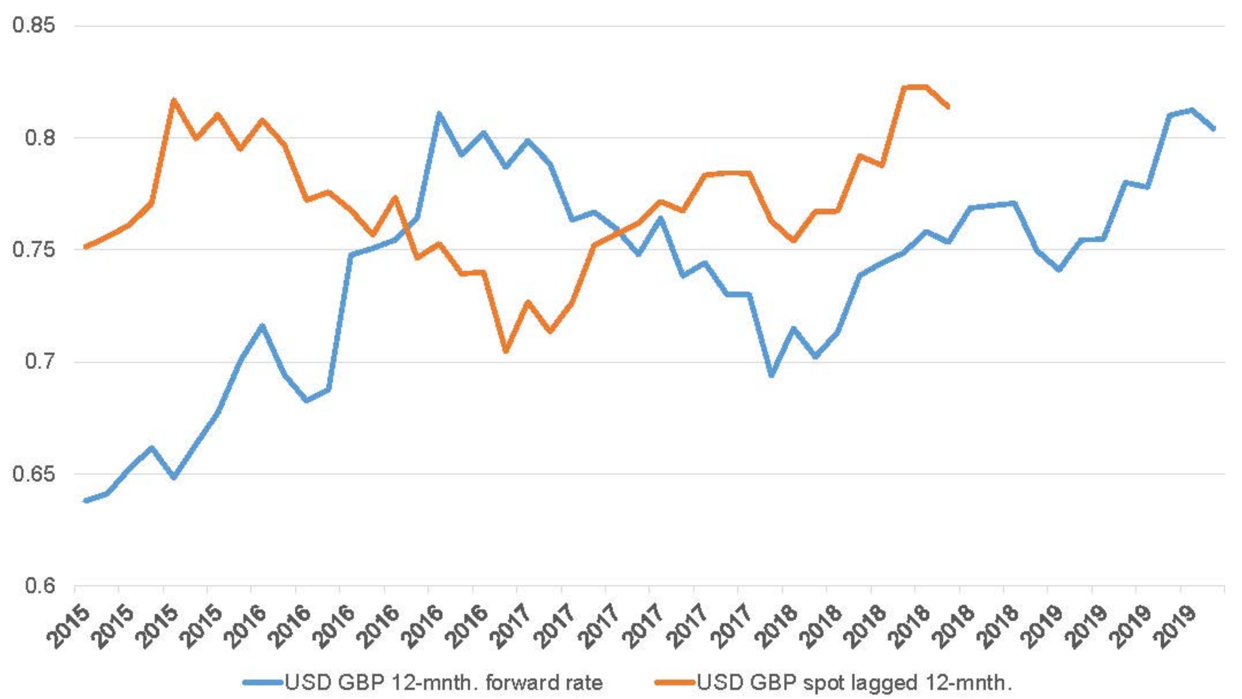

The U.S. dollar/British pound (USD/GBP) market quoted forward rate compared to the realized spot rate 12 months in the future. Prepared by Conning, Inc. Source: ©2020 Bloomberg, L.P.

Similar limitations might be observed regarding other forward rates, such as the forward yield curve, which is known to have limited forecasting power.

The use of econometric models for forecasting economic variables has gained some popularity over the last few years, not least within central banks, and may be a useful tool as well. Models such as the dynamic Nelson–Siegel model (Diebold and Li 2006) have been shown to have reasonable forecasting power for interest rates over short time horizons and have the advantage of offering a well-defined method for setting a benchmark mean in a reasonably automated fashion. A potential limitation is that the model parameters must be estimated in some way, and that is typically achieved using historical data. As such, some of the limitations are similar to those present when using pure historical data analysis. In particular, the forecasts from econometric models may be dependent on the length of the time series used in the estimation process, and one must make some expert judgments regarding which data are relevant.

Analyst views or the view of an investment committee may be an option for larger companies with sizable in-house analyst and investment teams. It may also be sensible and convenient for investment-, risk-, and capital-modeling systems to all be based on the same underlying assumptions. Indeed, some regulatory regimes actively require the use of the capital model in business decision-making (e.g., the Solvency II use test). Great care, however, needs to be taken to ensure that views are suitably conservative, and that observer bias does not overly impact the derived benchmarks. Investment teams will also often only have a view on expected return and not on risk metrics such as volatility and value at risk. In this case, some thought needs to be given to the implications of having one team define the best-estimate mean and another team the risk benchmarks.

9.6 Estimation of Risk Benchmarks—Volatility and Tail Effects

For regulatory capital purposes, the risk metric the user usually focuses on is the quality of the tails—in other words, the frequency and severity of extreme events. For other applications, one might be concerned with the standard deviation or a combination of standard deviation and tail metrics.

As with the estimation of the benchmark mean, several approaches might be considered when estimating targets for the standard deviation and tail metrics:

- **Historical data analysis:** This is the use of historical data to derive a suitable target volatility or tail metric.
- **Use of market-implied volatilities:** Here one uses the market prices of derivatives to imply a short-term volatility in a model.
- **Official regulatory targets:** In this approach we use values or methodologies defined within regulations as a basis for setting targets.

We can use historical data in several ways to determine appropriate risk characteristics for a model. Even for short-term calibrations, the use of long histories of data is not just preferred but likely to be necessary to adequately estimate benchmarks for standard deviations and tails. To estimate the standard deviation target from historical data the following two methods might be considered:

Data-window based. The user chooses a particular data window and the measure of the standard deviation to be used as a target. This requires some expert judgment as to which data window is most appropriate. For instance, U.S. interest rate volatility in the decade following the 2008 crisis was significantly lower than in previous decades, primarily due to central bank intervention. It may, however, not be prudent to assume that volatility will remain low even in the short term.

Fitted distribution based. In some circumstances it may be appropriate to first fit a known parametric distribution (e.g., normal or beta distribution) to the data and use the standard deviation of the fitted distribution as the basis of the benchmark. This technique may be particularly useful where historical data are bimodally distributed or contain one or more discontinuities in the distribution. This is because the measured standard deviation of a bimodal distribution is likely to be unrealistically high relative to an assumed continuous distribution of the type that most models produce.

As with the standard deviation, any benchmark for tail events is naturally best estimated using long histories of data; this is because tail events by their very nature are rare events. One must also decide whether an explicit quantitatively defined target should be set or if more qualitative acceptance criteria would be more appropriate.

If an explicit quantitative target is required or desired, then the first question is on which percentile we should base this—0.5th, 99.5th, 1st, or other? With this in place, one must then decide what value is reasonable. Using historical data for this task is not without limitations; as we have already stated, extreme events occur only rarely, typically with a frequency of a decade or more. There is no way of gaining any statistical certainty that the worst event or events observed in history equate to a particular percentile (e.g., the 1-in-100-year event). Even if we had more than 100 years of data, we would need to exercise some caution with respect to the relevance we place on an extreme event that happened more than 100 years ago.

We would be remiss not to mention that various techniques in the field of extreme value theory (EVT) have gained some prominence in the last few years. These may be a useful tool and provide additional data points but have several practical limitations. First, the application of EVT requires suitably expert personnel to apply the techniques and interpret the results. Second, the results of EVT are also dependent on data quantity and quality, as well as the assumptions that enter the algorithms used. As with all modeling tasks, the user must consider whether the added complexity and assumptions add or detract from the robustness of the estimate and the overall aims of the task. Clearly, though, if we are to use historical data, then some significant expert judgment is required to quantitatively express a view of the tails of distributions from the data.

Using market-implied volatilities of instruments to set the short-term risk distributions is another method commonly discussed. This method relies on first calibrating an arbitrage-free continuous-time stochastic model to the market price or implied volatility of derivatives at the valuation date.

Any risk premia within the model are then set to effect a change of measure to the real world and usually to bring the mean of the simulated distribution in line with the benchmark for the best estimate. In purely practical terms, calibrating a real-world model in this way requires a great deal of expertise. During extreme market events, the implied volatility of derivatives may also increase to levels that imply real-world distributions we are unlikely to ever observe. This creates calibrations that can be very unstable over time, leading to requirements for extra risk capital that might not be needed or justified. Also, this technique needs to rely on a deep and liquid derivatives market, which is not in existence for many asset classes and many economies. For these reasons, this technique is seldom used as the basis of determining short-term (or any other) risk distributions for real-world applications.

The final possibility we discuss here is the use of regulatory views of extreme risks as the basis for setting tail behavior benchmarks. This, of course, relies on a given regulator publishing or expressing some view, such as in the case of the Solvency II standard formula. However, regulators typically do not want analysts to regard these regulatory guidelines as blueprints for modeling, and they sometimes actively discourage users of stochastic models from using them in this way. For example, the standard model for Solvency II does not have downside risk for negative interest rates. Therefore, it may be best to regard these as a useful point of comparison in the validation process rather than de facto targets.

While setting the standard deviation or volatility targets may be relatively easy, setting tail targets is clearly very challenging. A more qualitative approach might be considered in order to mitigate some of the difficulties. For example, we may define acceptance criteria for a model calibration that place greater quantitative emphasis on the mean and standard deviation benchmarks but stipulate that "the model distribution should span all observed events in the last 50 years of data with 1,000 simulated paths." These types of qualitative judgments are easier to apply and allow greater flexibility when modeling limitations are encountered. At the same time, this stipulation to span all historical observed events with 1,000 simulated paths can be challenging. For example, in considering relative changes in interest rates, how do we deal with interest rates close to zero? If considering absolute changes, how do we deal with 300–basis point annual changes in the first half of the 1980s?

9.7 Model Implementation and Usage Considerations

The setting of model benchmarks would ideally also take into consideration the features and limitations of the available models. For instance, there is little advantage in setting a benchmark that implies that a distribution has high excess kurtosis and fat tails if the models available to the modeling team are capable of producing only normal distributions. Therefore, understanding the model and what one can reasonably expect of it in the context of matching short-term benchmarks is a first step prior to implementation.

The requirement to match initial, or time-zero, values for a given market while simultaneously matching particular benchmarks at a short time horizon makes the implementation more difficult than a longer-time-horizon calibration. In the case of long-term calibrations, the effects of changes in the initial conditions are usually minimal at the distant time horizon. Therefore, unless the long-term benchmarks change, one set of model parameters remains appropriate, often for many years. In the short-term case, the distributions of a simulated model will almost invariably be affected if the initial conditions change. This is illustrated in Figure 9.3, which shows the mean reversion behavior of two identical Ornstein–Uhlenbeck processes (i.e., the parameters are the same for both models). In this case the only difference is that model 1 is initialized from a time-zero value of 0.05 and model 2 from a value of 0.01. We see in the upper plot that both models mean-revert toward an identical value over time, as expected. The lower plot, which shows only the first three years of the simulation, illustrates that even after three years the effect of the initial condition has yet to wear off.

Figure 9.3. Mean reversion behavior of two identical Ornstein–Uhlenbeck processes initialized from different starting points

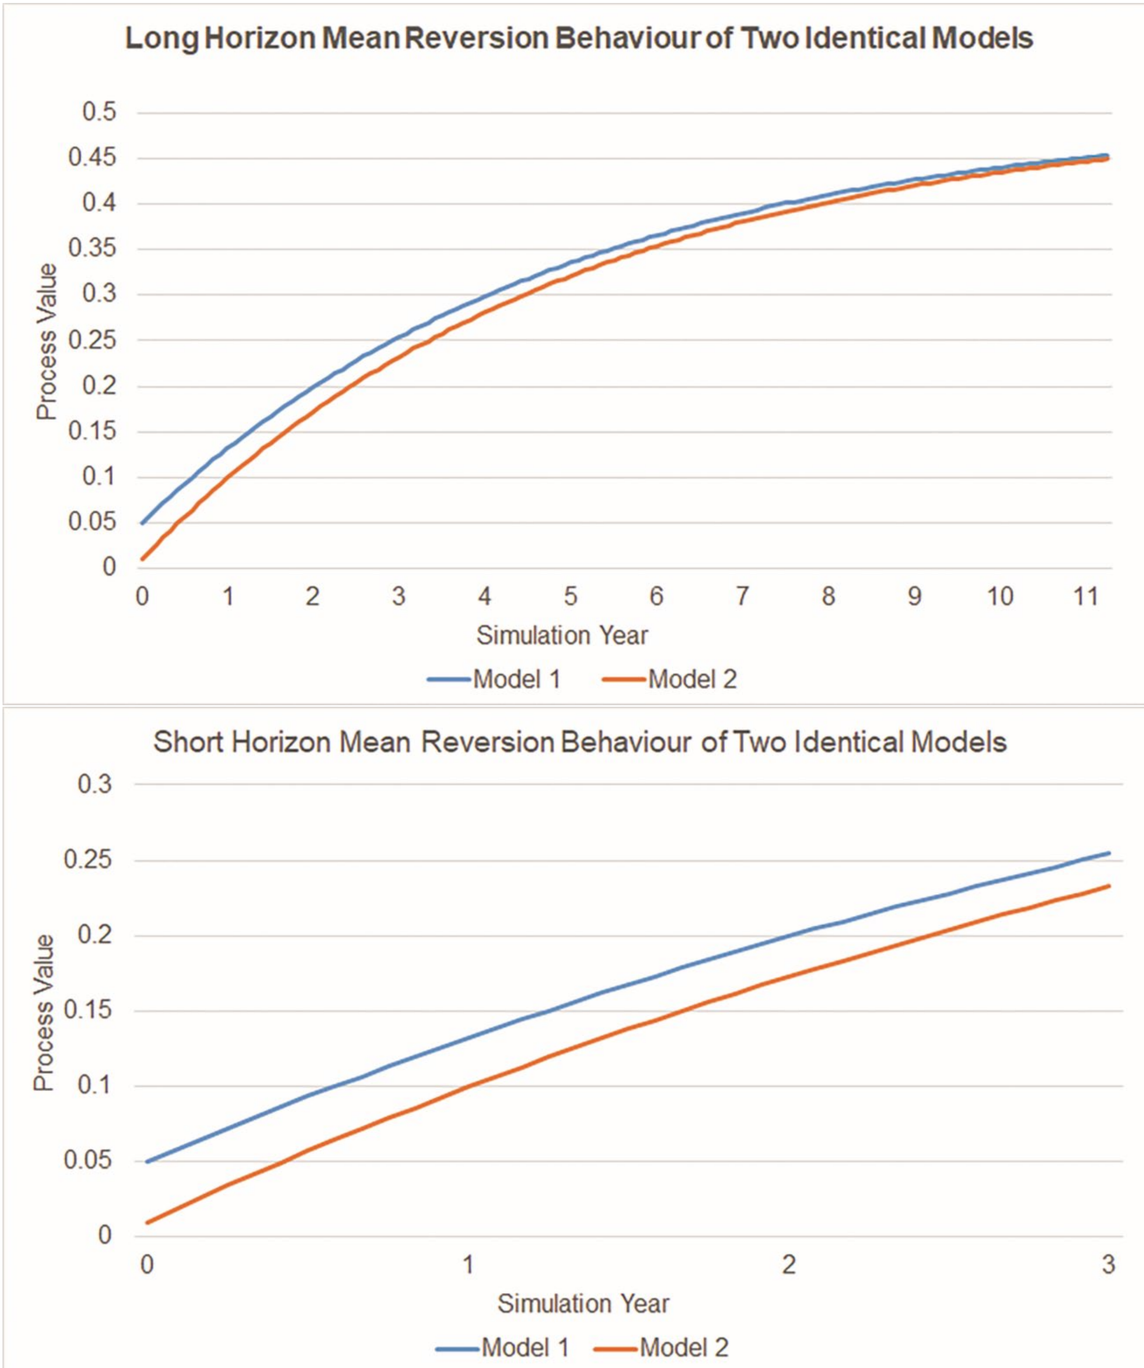

Prepared by Conning, Inc. Source: GEMS® Economic Scenario Generator scenario.

In practical terms, this would mean that if our short-term benchmark for the best estimate or mean of the distribution was the same in both cases but the initial conditions were different, we would need to re-estimate the parameters of the model to make them equivalent at the short-term horizon (e.g., by increasing the mean reversion speed of model 2).

In most cases, not just the simulated mean will change with the initial conditions, but all statistical properties of the model. This means that all the parameters of a model may very likely need to be intermittently re-estimated in order to keep the model in line with the specified benchmarks. Therefore, users of short-term calibrations should consider how much of this "model drift" is acceptable between valuation dates. This can be thought of as weighing up the need for a calibration to be in line with given benchmarks or targets and the desire to have parameter stability (i.e., that the model parameters do not change between valuation dates).

The propensity for models to drift in this way (or, more accurately, for markets to be volatile) places greater emphasis on model and system selection for short-term calibrations than it does for the long-term case.

The following are factors that practitioners must take into account when considering an ESG for the purposes of these types of calibration:

- **Flexibility of models:** Is (are) the available model (models) for a given asset class flexible enough to accommodate a wide range of different views at a short simulation time horizon?
- **Calibration tools:** Does the system come with inbuilt calibration tools that allow the user to adjust the model easily when initial conditions change? Do the calibration tools allow the user to calibrate to tail metrics as well as mean and standard deviation?
- **Extensibility of system:** Is it possible to extend models or augment deterministic scenarios into the system to ensure that particular types of events are accommodated in the short-term capital model?
- **Personnel:** To what extent do the models and system require expert personnel to implement and maintain the short-term calibrations?

It is also worth mentioning that a short-term calibration is usually restricted in its use to short time horizons. In some cases, a practitioner might create a calibration that produces reasonable dynamics at both long and short time horizons, but usually some compromises must be made. For instance, in order to match the desired long-term benchmarks, we may have to accept that the short-term benchmarks are achieved less accurately. It is, however, difficult to generalize, because the quality of a multi-horizon calibration of this kind depends on many factors.

9.8 Summary

In some circumstances an ESG user needs a model designed and calibrated specifically for the projection of assets and liabilities at short simulation horizons. These applications include regulatory capital calculations as well as capital models used for general risk management and investment decision-making.

In this chapter, we introduce the most important questions to consider when implementing an ESG for short-time-horizon simulations. There is no consensus on how a short-term calibration should be defined and implemented. We introduce the most commonly seen approaches and discuss in some detail the relative merits and limitations of each. The reader is advised that the final decision will depend on the application, but also on the models and systems available to him or her. Due to the likely requirement for ongoing maintenance and recalibration of the ESG, the expertise and availability of personnel with particular skill sets may also influence the details of the process chosen.

The practitioner needs to consider three areas at the outset: the establishment of suitable benchmarks, the task of finding appropriate model parameters that satisfy those benchmarks, and maintenance to ensure that the benchmarks and models remain appropriate.

Establishing short-term calibration benchmarks is arguably more difficult than for the longer-time-horizon case, because there is significant variability observed in the distributions and distributional properties of financial variables over different short time horizons. Two considerations are important when establishing calibration benchmarks: which risk metrics should be the focus of the target-setting procedure, and what methodology should one use to set the numerical value of the benchmark statistic?

Most applications will require calibration benchmarks for all or some of the following:

- *Initial conditions* (often called time-zero values), which act as a starting point for all simulations
- A *best estimate,* which most typically corresponds to the mean or median of the simulated distribution
- A *volatility benchmark* giving a target for the standard deviation of the simulated distribution
- *Relevant extremes* that the simulated distribution should exhibit

The setting of model benchmarks would ideally also take into consideration the features and limitations of the available models. The propensity for markets to be volatile places greater emphasis on model and system selection for short-term calibrations than it does for the long-term case. When selecting an ESG for the purposes of short-term calibration, practitioners should consider the model's flexibility, sufficiency of calibration tools, and extensibility to augment deterministic scenarios and the level of expertise required to maintain the calibrations.

References

Diebold, F. X., and C. Li, "Forecasting the Term Structure of Government Bond Yields," *Journal of Econometrics* 130, 2006, pp. 337–364.

Feldhütter, P., "Can Affine Models Match the Moments in Bond Yields?" *Quarterly Journal of Finance* 6:2, 2016.

Chapter 10: Software for Economic Scenario Generation

In the previous chapters we discussed the concept of an ESG and the important aspects of economic scenario generation, ranging from the development and maintenance of an ESG to the attributes of a good ESG. We are now interested in the situation where someone understands those issues and is in the position of needing an ESG for risk management. The purpose of this chapter is to address the following questions:

- What is economic scenario generation software (i.e., an ESG platform)?
- What features should an ESG platform have?
- What are the options for implementing ESG software (build, buy, or open source)?

We offer a few comments on the build-versus-buy decision from the last question, and then focus on open-source and commercial vendor solutions. Building your own ESG can be a satisfying and enlightening experience, as it permits a complete in-house understanding of your ESG systems and provides an opportunity to gear the ESG to the specific needs of your company—and it gets the various parties and stakeholders in the ESG process talking with one another at an early stage. However, building an ESG also requires constant maintenance and the need to develop the calibration, validation, and model-checking processes that we discuss in Chapter 4. Additionally, you will need your own (or access to other) data sources for initializing simulations to market conditions and for validation and model research. With all of these elements and responsibilities, building and operating your own ESG is expensive and requires considerable resources and development time. Insurance companies with access to staff having strong financial modeling talent, communication skills, and adequate budgets may find it not only feasible but possibly even a net benefit to build an ESG themselves.

We begin the chapter by discussing what is meant by economic scenario generation software. Some of the key features or capabilities that one might want the software to have are reviewed. We then discuss options for getting ESG software in place.

10.1 Economic Scenario Generation Software

Recall that an ESG is a computer-based model of an economic environment that is used to produce simulations of the joint behavior of financial market values and economic variables. Economic scenario generation software is an integrated suite of computer programs and graphical interfaces that permits the user to efficiently manage and generate economic scenarios. ESG software is the embodiment of the theory, methods, optimizations, calculations, and data management necessary for creating and applying simulated economic scenarios.

The developer of ESG software faces many considerations in choosing a programming language in which to develop the platform. Things are simpler for the user; indeed, from the user perspective the programming language matters mostly to the extent that it can deliver speed, security, and ease of use.

Features that may be part of economic scenario generation software include

- mechanisms for preparing and populating market data (i.e., initializing the models, stripping algorithms to obtain zero-coupon data, tools for estimating missing or lagged data such as GDP, inflation, or unemployment);
- tools for ensuring that the model starts the simulation with model yield curves that provide a close fit to the initial market yield curves;
- tools for calibrating component models to calibration targets;
- tools for querying ESG output and producing statistical or graphical summaries of the queried data;
- validation reports and methods for creating additional validation checks or reports;
- structures for creating user-defined models or discount frameworks (creating new financial variables or asset classes, defining a discount structure for pricing liabilities, creating new total-return series);
- an option to model at the security level (security-level cash flows, returns constructed from an actual security universe);
- tools for selecting scenarios with prescribed characteristics (i.e., path selection tool);
- systems for archiving and managing calibration parameters and simulation results; and
- the ability to custom-select economies and component models to run.

ESG software can range from simple systems that generate core economic series, like interest rates and inflation, to highly sophisticated systems that generate multiple economies with many asset classes and robust recalibration tools. All other things being equal, the simplest software that fits your particular purposes is often the best solution. However, the reality is that the uses of an ESG within an insurance organization evolve over time, and it is also necessary to anticipate—or at least have the flexibility to handle—important requirements that may emerge in the future. Once an ESG is adopted and integrated into an organization, it is costly to change, both in economic and in human-resources terms. This is particularly true when regulatory model approval is a factor. Therefore, an insurer may want to consider adopting an ESG software platform that has features that exceed its immediate usage needs and is capable of meeting anticipated future needs as well.

10.2 Open-Source ESG Software

By "open-source ESG software," we mean free downloadable ESG software that is fully documented and for which the models and production code are openly accessible to the user. Popular programming languages such as R and Python offer a broad range of libraries that can simulate many financial and economic variables. However, ESG software requires that the individual modules be engineered as components of the whole and integrated into a single simulation platform. Consequently, ad hoc software libraries, while very useful, do not qualify as ESG software.

The most widely known open-source ESG software is the American Academy of Actuaries (AAA) model. A second, well-known open-source ESG software is the Casualty Actuarial Society/Society of Actuaries model (i.e., the Ahlgrim–D'Arcy–Gorvett, or ADG, model) mentioned in prior chapters. We will discuss both of these models.

10.2.1 The American Academy of Actuaries Model

You may access the AAA model—and download the simulation engine and related documentation—from either of the following websites:

- https://www.actuary.org/content/economic-scenario-generators
- https://www.soa.org/resources/tables-calcs-tools/research-scenario/

According to the Society of Actuaries website: "The American Academy of Actuaries and the Society of Actuaries (SOA) have joined resources to manage the economic scenario generators used in regulatory reserve and capital calculations.

"The SOA will provide frontline support and maintenance for the economic scenario generators. A joint Society of Actuaries/American Academy of Actuaries oversight group will oversee the generators and assist the SOA in providing technical support and direction for the current and future versions of the generators."

The model is spreadsheet-based and includes the following variables:

- U.S. Treasury interest rates (i.e., risk-free interest rates)
- Money market bond index fund return
- Intermediate-term government bond index fund return
- Long-term corporate bond index fund return
- Domestic diversified equity index fund return
- International diversified equity index fund return
- Intermediate risk equity index fund return
- Aggressive/exotic equity index fund return

Bond fund returns are derived from the simulated Treasury interest rates, a duration approximation, and a random error term.

Figure 10.1 shows output from the AAA model, run as of June 30, 2019, using the default calibration parameters provided with the simulation engine. The histogram and fan chart show the distribution of the three-month Treasury interest rate at the simulated 10-year horizon. The starting/initial value of the three-month Treasury interest rate is 2.12%.

Figure 10.1. AAA model output

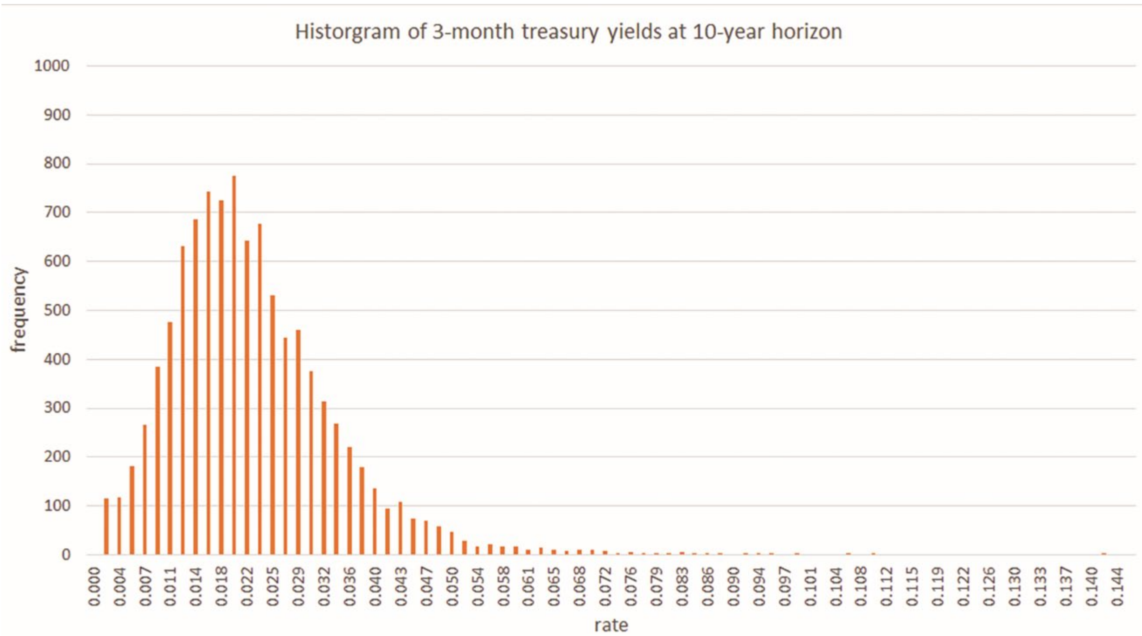

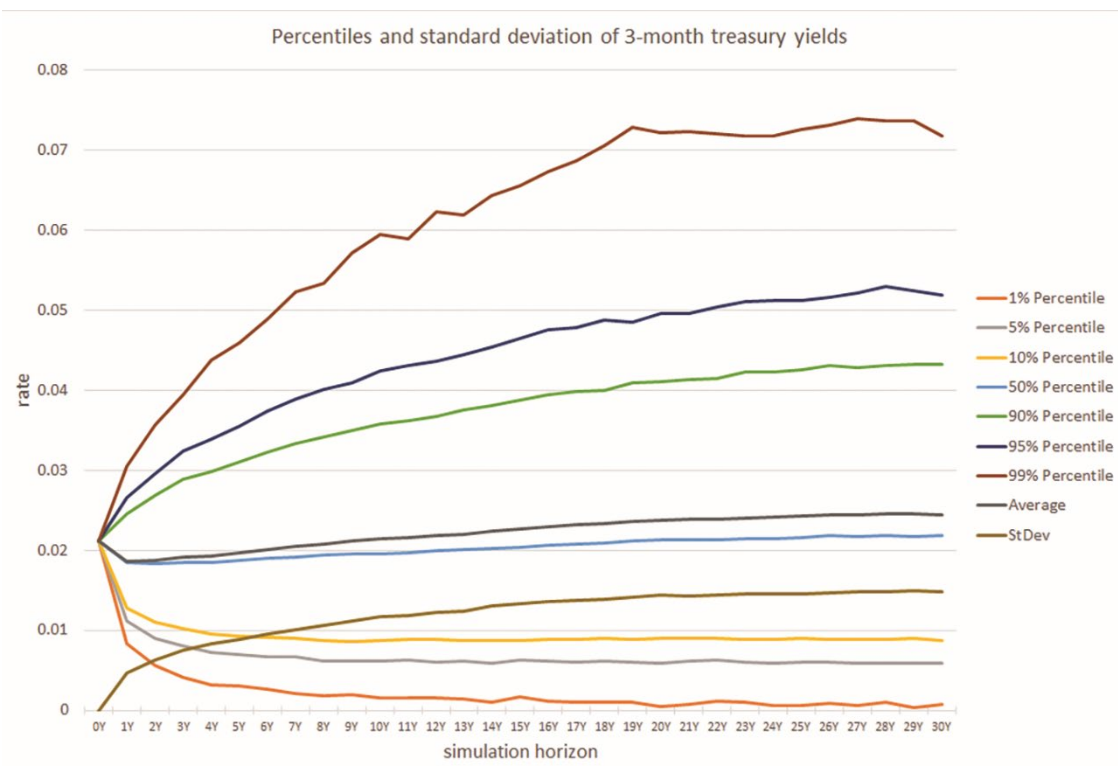

Prepared by Conning, Inc.

The distribution of the 10-year Treasury interest rate at the 10-year simulation horizon is shown in Figure 10.2. The starting/initial value of the 10-year Treasury interest rate is 2.00%.

Figure 10.2. AAA model output

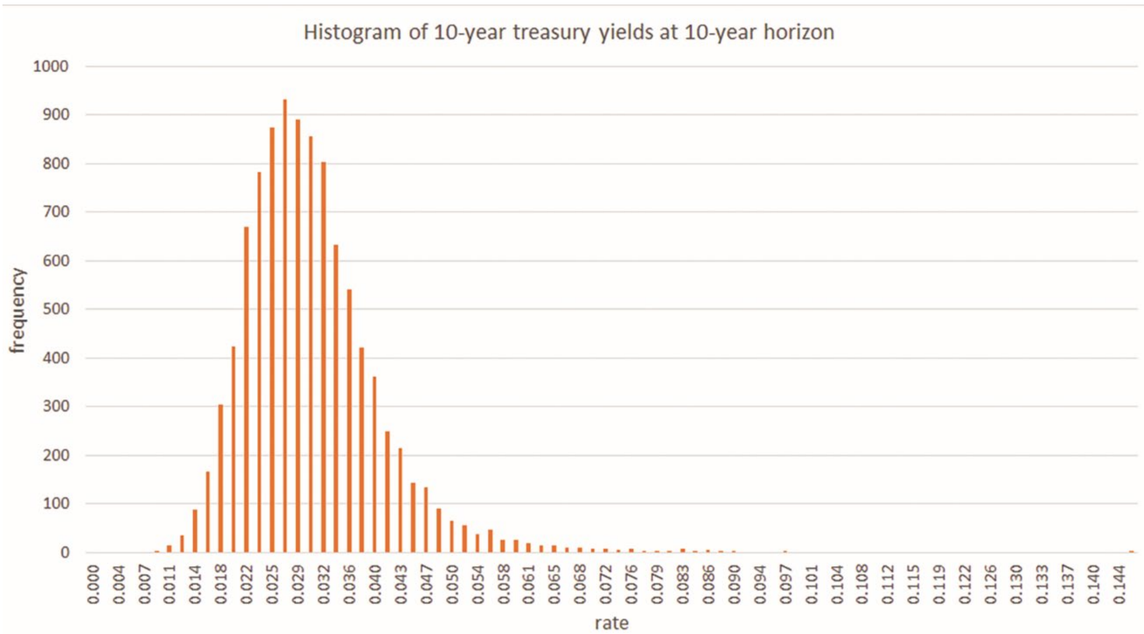

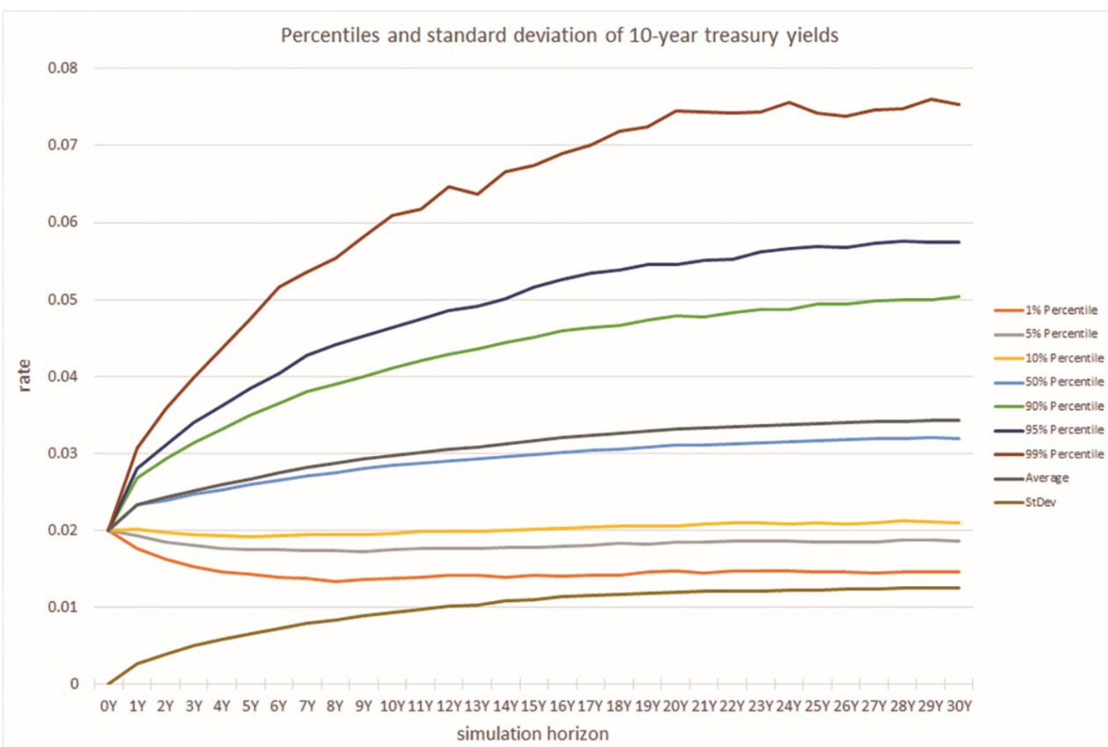

Considering that these distributions are measured at the 10-year simulation horizon, the range of the distribution of interest rates for the default calibration might be too narrow for some real-world applications. A user of the AAA model has direct control over the model parameters through the "Parameters" tab but does not have a recalibration tool for adjusting the model to specific calibration requirements.

The AAA model is widely used for regulatory compliance work.

10.2.2 Casualty Actuarial Society/Society of Actuaries (ADG) Model

The Casualty Actuarial Society/Society of Actuaries (CAS/SOA) model was developed by Kevin Ahlgrim, Stephen D'Arcy, and Rick Gorvett (2004, 2005), and is discussed in more depth in Chapters 6 and 7. It is available online, together with supporting documentation and supplemental research, at

- https://www.casact.org/research/econ/.

The model is spreadsheet-based and includes the following variables:

- Inflation
- Real interest rates
- Nominal interest rates, which are implied from the processes for inflation and real interest rates
- Large- and small-stock returns
- Equity dividends
- Returns from real estate investments
- Rate of unemployment

The same researchers also provide a comparison of the AAA model and the CAS/SOA model in the paper "A Comparison of Actuarial Financial Scenario Generators" (Ahlgrim, D'Arcy, and Gorvett 2008), available at

- https://www.variancejournal.org/issues/?fa=article&abstrID=6467.

The CAS/SOA model offers different economic variables than the AAA model, based on different stochastic dynamics.

10.3 Commercial Vendor ESG Products

The user can choose from many commercially available ESG software products. ESG software vendors include

- Moody’s (previously Barrie & Hibbert);
- Conning (GEMS);
- Ortec Finance;
- Willis Towers Watson (STAR ESG); and
- Numerix.

ESG software products from commercial vendors typically incorporate features and services that are convenient and valuable to the user, some of which are

- official calibration parameter releases and associated validation reports;
- suite of recalibration tools that permit the user to adjust calibration properties;
- cloud-based solutions that lever simulation speed and can allow easier management of information technology;
- ESG software that may be integrated into a widely used actuarial liability software platform;
- services for implementing ESG models and ongoing model maintenance; and
- services for assisting an insurance company in regulatory approval processes.

Most commercial vendors offer stand-alone ESG scenario file services. This has the advantage of reduced cost and avoids the need to install and manage ESG software and can be a viable solution for simpler use cases.

10.4 Summary

Users of economic scenario generators have a wide choice of ESG software solutions. Open-source ESGs, with relatively limited functionality, are available. A wide range of ESG software solutions are available from several commercial vendors. Commercial ESG software is widely available in the cloud, which can offer speed and cost-saving advantages.

The company use case, anticipated needs, and information technology requirements should drive the choice of the ESG software solution. A solid understanding of both the issues involved in the development and maintenance of an ESG and the attributes of a good ESG is an important guidepost in selecting the appropriate ESG software.

References

Ahlgrim, K. C., S. P. D’Arcy, and R. W. Gorvett, “Report on Modeling of Economic Series Coordinated with Interest Rate Scenarios,” Arlington, VA: Casualty Actuarial Society, 2004, https://www.casact.org/research/econ/.

Ahlgrim, K. C., S. P. D'Arcy, and R. W. Gorvett, "Modeling Financial Scenarios: A Framework for the Actuarial Profession," *Proceedings of the Casualty Actuarial Society* 92, 2005, pp. 177–238.

Ahlgrim, K. C., S. P. D'Arcy, and R. W. Gorvett, "A Comparison of Actuarial Financial Scenario Generators," *Variance* 2:1, 2008, pp. 111–134.

Chapter 11: Guide to the Literature on Economic Scenario Generation

In vibrant research areas such as those related to economic scenario generators, it is difficult to thoroughly survey and identify the literature, since the pool of relevant and potentially meaningful papers and research reports is ever growing. The best one can do is provide an annotated bibliography of the particularly significant or "classic" papers, a list of additional papers with which researchers may want to become familiar, and a few suggestions for where future literature of relevance may be found. These are our goals in this chapter.

Section 11.1 identifies the characteristics of the ESG literature as a whole, and the contexts within which potentially useful material may be found. Section 11.2 provides descriptions of several of the major, "classic" papers that deal directly with ESGs. These papers include material with which every designer, builder, and probably even user of ESGs should be familiar.

Section 11.3 also provides an annotated bibliography—of major papers categorized according to their particular subject matter. Deeper understanding of this material is appropriate for those involved or interested in the specific modules or aspects of an ESG.

Section 11.4 comprises a list, without comments, of additional readings in ESG-related areas—sources that provide either deeper analysis of material or alternative approaches to modeling. The final section, Section 11.5, offers brief suggestions for where future research may likely be found as it emerges.

In general, within each category, we have listed papers in chronological order of publication.

11.1 Characteristics of the Published Literature Relating to Economic Scenario Generation

Actuarial science as a whole is one of the most interdisciplinary (or, if you prefer, multidisciplinary) of professions. It is impossible to adequately analyze data, or what that data may mean for the future, without knowing something about where they come from, the circumstance under which the data emerged, the nature of changes between historical and prospective environments, etc. Thus, any subject matter associated with the socioeconomic, financial, technological, cultural, political, etc. environment is potentially relevant to the evaluation and quantification of risk.

Similarly, building and using ESGs involves a multidisciplinary context. Indeed, while it is critical for an actuary designing an ESG to keep in mind the purpose and function that the ESG will serve in the context of the organization, the specific issues entering into an ESG model are areas of both academic and institutional scholarship and research in themselves: economics, finance, econometrics, and statistics. Thus, much of the research and guidance that one needs to build and maintain an ESG can be found in the literature streams associated with those topics.

There are at least two (although probably more) ways of categorizing papers and research relevant to economic scenario generation. One approach is to look at the ESG process based on the specific modeling function—that is, to classify papers according to which function associated with building a model is relevant, such as, for example,

- theory associated with model design;
- estimation of model parameters; or
- validation of the model.

This is a perfectly reasonable way to categorize the literature, and each of the three functions shown above involves different considerations and techniques.

The second approach to categorizing papers—and essentially the one we employ in this chapter—is by the type of economic variable being generated, such as, for example,

- interest rates;
- equity returns; or
- real estate investment factors.

This is probably the more convenient approach to categorization, since most research is published in journals that belong to a particular domain, or subject matter area.

In Sections 11.3 and 11.4, we review some of the relevant literature according to the following categorization scheme:

- Interest rates and credit risk
- Equity pricing and modeling
- Other literature

First, however, in Section 11.2, we summarize some of the classic literature specifically devoted to ESGs as a whole.

11.2 Discussion of Some Classic ESG Papers

Literature associated with designing, building, and using ESGs has been published over the last 30 years. While not exclusively so, much of this research work has been done by actuaries. Since much of the detail associated with ESGs is relevant to both life and property/casualty applications, we present below descriptions of some of the most cited and foundational of the publicly available ESG literature regardless of the specific actuarial discipline.

Wilkie, A. D., "A Stochastic Investment Model for Actuarial Use," *Transactions of the Faculty of Actuaries* 39, 1984, pp. 341–403.

Wilkie, A. D., "More on a Stochastic Model for Actuarial Use," *British Actuarial Journal* 1, 1996, pp. 777–964.

The Wilkie model, as introduced (1984) and then elaborated upon (1996), is one of the original and most cited examples of an economic scenario generator. It has served as a foundation, as well as a point of reference, for many later efforts to build ESGs. Although originally developed for life insurance applications, the framework for simulating economic scenarios is equally valid for property/casualty.

The Wilkie model uses a cascade structure, beginning with inflation as the initial or primary variable. For inflation, a first-order autoregressive process is used, and then other economic and financial variables are built upon the inflation simulation.

Wilkie uses and develops several sophisticated statistical tools—e.g., VAR (vector autoregression), ARCH (autoregressive conditional heteroscedasticity)—and estimates parameters based on data over a long period, from 1923 to 1994. However, he also uses a fair amount of judgment and personal preference to deal with issues that come up during his analysis.

The extensive list of variables modeled by Wilkie includes the following:

- Inflation—the primary (or driving force) variable; modeled as an AR(1) process
- Wages—tested for potential relationships with inflation
- Dividend yields on common stock—an AR(1) process with some dependence on inflation
- Dividends—based on current and past inflation, dividend yields, and dividends
- Short-term interest rates—AR(1) process for long-short spread
- Long-term interest rates—based on past inflation and the real interest rate
- Real estate—based on property yields and income
- Foreign exchange—fluctuations around foreign exchange rates implied by purchasing power parity

Tilley, J. A., "An Actuarial Layman's Guide to Building Stochastic Interest Rate Generators," *Transactions of the Society of Actuaries* 44, 1992, pp. 509–564.

This paper is a discussion of the theory of arbitrage-free pricing of interest-rate-contingent cash flows. Important concepts discussed include arbitrage, complete markets, expected value pricing in the risk-neutral world, and stochastic interest rate models. The paper is written for actuaries and appears in the SOA's scholarly publication, but it is descriptive and nonmathematical.

Here is a quote from the author's concluding remarks: "This paper has shown that it is possible to gain a practical understanding of key concepts in financial economics without having to resort to a study of the mathematics of stochastic processes. Based on the assumptions that opportunities for riskless arbitrage do not exist and that the financial markets are complete, it was shown that a theory could be developed for pricing interest-rate-contingent cash-flow streams relative to the prices of all zero-coupon bonds, which are taken as exogenous inputs to the theory. The arbitrage-free prices are calculated in a straightforward manner by the expected-present-value algorithm, a technique that lies at the heart of actuarial science" (p. 536).

The discussions appearing after the paper, and the author's responses, are also valuable contributions.

Hibbert, J., P. Mowbray, and C. Turnbull, "A Stochastic Asset Model & Calibration for Long-Term Financial Planning Purposes," Technical Report, Barrie & Hibbert Limited, 2001.

This extensive paper describes a model that generates consistent values across a large number of economic and financial variables, including the term structure of interest rates (real and nominal), inflation rates, equity returns, and dividends. The paper reviews the history of interest rates, inflation rates, and equity returns over the last 100 years, as well as for more recent periods.

The authors employ a two-factor Hull–White model for the real interest rate. The short-term rate is mean-reverting, where the long-run mean itself is a stochastic process. Similarly, inflation is also a two-factor model, with a double-mean-reversion process.

Equity returns are modeled by an equity premium, or excess return over and above the nominal interest rate, employing a Markov regime-switching model with one regime having a higher expected return and lower variance and the other regime a slightly lower (even negative) expected return but much larger variance.

The equity dividend yield model is a one-factor first-order autoregressive process. The paper provides a significant description of the calibration process for this model. Two sets of parameter values are illustrated, one that allows negative interest rates and does not include a risk premium for long-term interest rates, and another that limits nominal interest rates to positive values. The model simulates scenarios over a 30-year horizon based on monthly steps. The resulting means, standard deviations, and distributions are then compared to illustrate the impact of this change in calibration.

Finally, the paper compares the model results to the output from the Wilkie model. The Wilkie model is shown to generate inconsistent relationships among inflation, bank interest rates, and the yield on consols. The autoregressive feature of equity returns included in the Wilkie model generates a distribution over a long-term horizon that is much more compact than historical experience would indicate.

Ahlgrim, K. C., S. P. D'Arcy, and R. W. Gorvett, "Report on Modeling of Economic Series Coordinated with Interest Rate Scenarios," Arlington, VA: Casualty Actuarial Society, 2004, https://www.casact.org/research/econ/.

Ahlgrim, K. C., S. P. D'Arcy, and R. W. Gorvett, "Modeling Financial Scenarios: A Framework for the Actuarial Profession," *Proceedings of the Casualty Actuarial Society* 92, 2005, pp. 177–238, https://www.casact.org/pubs/proceed/proceed05/05187.pdf.

> In May 2001, the CAS and the SOA jointly issued a request for proposals (RFP) on the research topic "Modeling of Economic Series Coordinated with Interest Rate Scenarios." The 2004 report and the 2005 paper summarize the authors' research project, initiated in response to the joint RFP. The authors constructed a financial scenario model that simulates a variety of economic variables—including U.S. interest rates, inflation, equity returns, dividend yields, real estate returns, and unemployment—over a 50-year period.
>
> The report and paper discuss modeling issues, review the literature, and describe the components of the model. The authors also discuss data and calibration, as well as simulated output. The paper serves as a practical guide not just to this particular model but to designing, building, and running an ESG in general. Uses for the model, as noted by the authors, include dynamic financial analysis, solvency margin determination, cash flow testing, operational planning, and other insurer financial analyses.

Ahlgrim, K. C, S. P. D'Arcy, and R. W. Gorvett, "A Comparison of Actuarial Financial Scenario Generators," *Variance: Advancing the Science of Risk* 2:1, 2008, pp. 111–134, http://www.variancejournal.org/issues/?fa=article&abstrID=6467.

> This paper is associated with the authors' research project, funded by the CAS and SOA, on "Modeling of Economic Series Coordinated with Interest Rate Scenarios." They discuss and compare the CAS/SOA model with the American Academy of Actuaries (AAA) model for equity returns and interest rates; the latter model supports the C3 Phase II variable annuity risk-based capital requirements. The comparisons are used to determine the impact of the use of different assumptions and parameter selection on the modeling process.
>
> The paper's comparison involves two examples of cash flow testing. The first example is based on a single premium, 10-year term life insurance policy, with the net single premium invested in T-bills (50%), large-cap stocks (25%), and small-cap stocks (25%). The second example is based on a $10 million property-liability loss reserve situation. This reserve is presumed to be supported by $10 million in assets, because property-liability insurance accounting does not permit discounting of loss reserves. The CAS/SOA model generally leads to a wider set of distributions than the AAA model.

11.3 Annotated Bibliography of Some Recent Papers by Category

Below are brief comments on some of the important papers in each of the three categories mentioned in Section 11.1.

11.3.1 Interest Rates and Credit Risk

Vasicek, O., "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics* 5, 1977, pp. 177–188.

The Vasicek model is often used as a simple example of a general equilibrium interest rate model. It is similar to the Cox–Ingersoll–Ross model, except that the volatility of the process is constant and does not depend upon the level of the interest rate. The Vasicek model for the evolution of the short-term interest rate is

$$dr_t = \kappa(\theta - r_t)dt + \sigma dB_t,$$

where r is the short-term (instantaneous) rate of interest, κ is the speed of mean reversion, θ is the long-run mean, σ is a volatility parameter, and dB is a standard Wiener process.

Cox, J. C., J. E. Ingersoll, and S. A. Ross, "A Theory of the Term Structure of Interest Rates," *Econometrica* 53:2, 1985, pp. 385–407.

This paper is one of the classics of the literature involving the term structure of interest rates. The CIR model is frequently used for academic studies of interest rates, as well as by practitioners. It is a relatively simple general equilibrium model of term structure, and it has appealing characteristics regarding the evolution of the short-term default-free interest rate, such as the following:

- There are no negative interest rates (when looked at in a continuous-time framework).
- The rate mean-reverts to a long-run value.
- The volatility is (square-root) related to the level of the interest rate.

The continuous version of the evolution of the short rate involves a single state variable, r, and has the form

$$dr_t = \kappa(\theta - r_t)dt + \sigma\sqrt{r_t}dB_t,$$

where r is the short-term (instantaneous) rate of interest (e.g., the return on U.S. Treasury bills), κ is the speed of mean reversion, θ is the long-run mean, σ is a volatility parameter, and dB is a standard Wiener process.

Pedersen, H. W., E. S. W. Shiu, and A. E. Thorlacius, "Arbitrage-Free Pricing of Interest Rate Contingent Claims," *Transactions of the Society of Actuaries* 41, 1989, pp. 231–279.

The authors use a discrete-time binomial lattice framework to investigate the pricing of bond options and interest-sensitive cash flows, with the result being an arbitrage-free binomial model that allows the initial term structure of interest rates to be prescribed exogenously. This allows the model price for each stream of fixed and certain cash flows to be set as the market price. Their approach can be viewed as a generalization of Ho and Lee (1986—see Section 11.4, "Papers for Further Reading").

Tilley, J. A., "A Stochastic Yield Curve for Asset/Liability Simulation," paper delivered at Actuarial Approach for Financial Risks: First AFIR International Colloquium, Paris, April 23–27, 1990, International Actuarial Association.

The paper presents an empirical approach to long-term asset/liability simulations. The model has dynamic U.S. Treasury yield curves but no other variables. It is intended for modeling U.S. life insurer or pension assets and liabilities.

The yields are based on the bid prices for on-the-run U.S. Treasury securities, from 1981 to 1989, with a yield range from three months to 30 years. To those data, orthogonal third-degree polynomials were fit; it turns out that the dynamics are described by an autoregressive process of order two. Tilley reports two interesting results:

- The fitted coefficients exhibit mean reversion, with the result that the yield curves tend to revert over the long term to a normal, positively sloped form centered near 8.3%.
- The residuals have fat-tailed distributions that are explainable by a mixture of two normal distributions.

Litterman, R., and J. A. Scheinkman, "Common Factors Affecting Bond Returns," *Journal of Fixed Income* 1:1, 1991, pp. 54–61.

The authors identified just three types of shifts—the level, steepness, and curvature of the term structure—that explain almost 99% of the historical observed variance in interest rates. In fact, the first factor alone—level shifts up or down in the term structure—explained almost 90% of the total variation in term structure movements over time.

Heath, D., R. Jarrow, and A. Morton, "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation," *Econometrica* 60:1, 1992, pp. 77–105.

This important and frequently cited paper provides a unifying framework for understanding arbitrage pricing theory associated with default-free bonds. The authors claim that previous models—such as Vasicek (1977), Brennan and Schwartz (1979), and Ho and Lee (1986—see Section 11.4, "Papers for Further Reading")—are special cases of their framework.

The HJM approach models the entire term structure—taking the initial forward rate curve as given—through a series of forward rates processes (as opposed to modeling the evolution of the short rate over time). The authors specify a general continuous-time stochastic process for the evolution of that curve across time in such a way that it has an equivalent martingale probability measure.

Chan, K. C., G. A. Karolyi, F. A. Longstaff, and A. B. Sanders, "An Empirical Comparison of Alternative Models of the Short-Term Interest Rate," *Journal of Finance* 47:3, 1992, pp. 1209–1227.

This paper provides estimates of parameters for a class of term structure models, based on monthly interest rate data from 1964 to 1989. Parameters are calculated using the generalized method of moments framework. Among other things, the authors find that the volatility of interest rates is very sensitive to the interest rate level.

Hull, J., and A. White, "One-Factor Interest Rate Models and the Valuation of Interest-Rate Derivative Securities," *Journal of Financial and Quantitative Analysis* 28:2, 1993, pp. 235–254.

This important paper, which introduces the well-known Hull–White model, discusses different approaches to arbitrage-free models of the term structure of interest rates. The authors' model involves a numerical procedure for constructing one-factor models that are Markovian and consistent with the initial term structure—a characteristic of some interest rate models that is highly valued for certain financial applications, such as in support of short-term trading.

One approach discussed in the paper is to specify a process for the short rate that has one or two unknown functions of the interest rate drift, and then estimate the functions to make the model fit the initial term structure. The authors cite several examples of this approach, including models by Vasicek (1977), Cox, Ingersoll, and Ross (1985), and Black and Karasinski (1991). The authors also describe a general procedure based on a trinomial tree, which can also be used to value options on bonds and interest rates.

Ahlgrim, K. C., S. P. D'Arcy, and R. W. Gorvett, "Parameterizing Interest Rate Models," *Casualty Actuarial Society Forum*, Summer, 1999, pp. 1–50, http://www.casact.org/pubs/forum/99sforum/99sf001.pdf.

The authors examine the statistical characteristics of historical interest rate movements from 1953 to 1999. They then look at several common interest rate models, identify some of their key characteristics, and discuss the selection of parameters of those models to assess their abilities to fit historical interest rate movements.

Chen, R.-R., and L. Scott, "Multi-Factor Cox–Ingersoll–Ross Models of the Term Structure: Estimates and Tests from a Kalman Filter Model," *Journal of Real Estate Finance and Economics* 27:2, 2003, 143–172.

The paper presents a method for estimating multifactor versions of the Cox, Ingersoll, and Ross (1985) model of the term structure of interest rates. The fixed parameters in one-, two-, and three-factor models are estimated by applying an approximate maximum likelihood estimator in a state-space model using data for the U.S. Treasury market. A nonlinear Kalman filter is used to estimate the unobservable factors.

The authors show that multifactor models are necessary to characterize the changing shape of the yield curve over time, and the statistical tests support the case for two- and three-factor models. They reject two-factor models in favor of three-factor models and argue that a three-factor model can incorporate random variation in short-term interest rates, long-term rates, and interest rate volatility.

Duffie, D., and K. J. Singleton, "Modeling Term Structure of Defaultable Bonds," *Review of Financial Studies* 12:4, 1999, 687–720.

This paper presents reduced-form models of the valuation of contingent claims subject to default risk, focusing on applications to the term structure of corporate and sovereign bonds. The default is treated as an unpredictable event governed by a hazard-rate process. The losses at default are specified in terms of the fractional reduction in market value when default occurs. Under this setting defaultable bonds can be priced as default-free bonds with the default-free discount rate replaced with the default-adjusted discount rate, and therefore all of the standard analysis of risk-free rate models can be applied.

The paper applies this framework to many existing models and various credit derivatives, and it shows that the pricing framework produces robust results.

Feldhütter, P., "Can Affine Models Match the Moments in Bond Yields?" *Quarterly Journal of Finance* 6:2, 2016.

Historically, the data suggest that U.S. interest rates vary over time with respect to their risk premia and volatility. This paper examines how well three-factor affine term structure models can capture the dynamic excess return and yield volatility (as well as higher moments) characteristics of bonds. The author finds that the three-factor models tested are less than fully successful in jointly modeling return predictability and yield volatility. The paper is a useful guide and companion to testing yield curve models, and for descriptions of some other models in the academic literature.

11.3.2 Equity Pricing and Modeling

Amaral, L. A. N., V. Plerou, P. Gopikrishnan, M. Meyer, and H. E. Stanley, "The Distribution of Returns of Stock Prices," *International Journal of Theoretical and Applied Finance* 3:3, 2000, pp. 365–369.

This is an empirical study of stock returns of individual companies. Two sources provide the data: the trades and quotes database for 1994–1995 and the Center for Research in

Security Prices database for 1962–1996. The distribution of returns is studied over various timescales, from several minutes to several years. There are three findings:

1. The distribution of normalized returns for individual companies is consistent with a power-law behavior characterized by an exponent $\alpha \approx 3$. (For perspective, this is outside the stable Lévy range of $0 < \alpha < 2$.)

2. The distributions of returns retain the same functional form—a power-law decay—for timescales from five minutes to 16 days.

3. For timescales greater than 16 days, the distribution of returns appears to slowly converge to a Gaussian distribution.

Hardy, M. R., "A Regime-Switching Model of Long-Term Stock Returns," *North American Actuarial Journal* 5:2, 2001, pp. 41–53.

The author uses a log-normal regime-switching model for equity returns. The model is parameterized using monthly S&P 500 and Toronto Stock Exchange data and is compared with other possible models of equity returns.

Durham, G. B., "Risk-Neutral Modeling with Affine and Nonaffine Models," *Journal of Financial Econometrics* 11:4, 2013, pp. 650–681.

This is an important development in asset-pricing models, although necessarily long and complex. The author uses option prices on stocks to obtain information about the future price dynamics of those stocks but recognizes an issue: the implied price dynamics of the underlying asset derived from options prices follow a risk-neutral measure while the price dynamics themselves follow a physical or real-world measure. By using both option and underlying asset price data, the article describes new techniques for joint analysis of the real-world and risk-neutral models.

The paper considers seven log volatility models and seven affine models classified as follows: no jumps, jumps in returns alone, and jumps in both returns and volatility. A typical model assessment is (i) derive the implied volatility from the observed option prices; (ii) compute the log-likelihood; and then (iii) examine diagnostics (e.g., looking for standard normal and independent residuals).

The author finds that "log volatility models perform dramatically better than affine models, but that some evidence of misspecification remains" (p. 650). Further, "all the models are rejected at far beyond conventional significance levels on at least one test. One would have to be almost hopelessly optimistic to believe that any of the models examined in this article was the true data-generating process. Failure to reject a model in exercises such as this should more likely be interpreted as a sign of insufficient sample size or tests with low power rather than an indication that one has found the true data-generating process. Powerful diagnostics are a good thing; failure to find evidence of defective models is a serious liability" (p. 679).

11.3.3 Other Literature

Heston, S. L., "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies* 6:2, 1993, pp. 327–343.

One of the most popular stochastic volatility models for equity prices and equity derivative pricing, the Heston model extends the Black and Scholes (1973) model by relaxing the constant volatility assumption. The model incorporates a stochastic process (CIR) for instantaneous variance. The model can accommodate important features such as non-log-normal distribution of equity returns, leverage-effect, and volatility clustering and still remain analytically tractable. The author derives closed-form solutions for equity options and illustrates the importance of (negative) correlation between volatility and equity price to explain skewness of equity returns.

Bates, D. S., "Jumps and Stochastic Volatility: Exchange Rate Process Implicit Deutsche Mark Options," *Review of Financial Studies* 9:1, 1996, pp. 69–107.

This paper presents a stochastic volatility jump diffusion model, the "Bates model." The Bates model extends the Heston (1993) model by adding a jump process to the price process and inherits many desirable features of the Heston model. But the Heston model often fails to generate extreme events in the equity markets properly. The jump process brings into the model the possibility of large price changes that cannot be generated by a continuous diffusion process. By adding a jump process, the Bates model can generate extreme events with a frequency and magnitude closer to that observed in the market. The Bates model is considered to be the most realistic continuous-time model for equity prices.

The author applies the model to exchange rate data, but it can be used readily for equity data. With properly estimated parameters, the Bates model can produce desirable features for the equity model such as fat tails in the return distribution, persistent volatilities, jump clustering, severe drawdown, and realistic interaction among multiple equity indices. The model also allows for an accurate semi-closed form option-pricing formula and efficient calibration procedures.

Wang, Y.-C., J.-L. Wu, and Y.-H. Lai, "A Revisit to the Dependence Structure between the Stock and Foreign Exchange Markets: A Dependence-Switching Copula Approach," *Journal of Banking and Finance* 37:5, 2013, pp. 1706–1719.

This paper has great import for an international ESG. It expounds a dependence-switching copula model to investigate the dependence structure between stock markets and foreign exchange markets. In particular, it looks at the four different combinations of market status: rising or falling stock values (i.e., bull and bear markets) combined with appreciating or depreciating currency values. The model uses daily stock returns and exchange rate changes for six major industrial countries (Canada, France, Germany, Italy, Japan, and the United Kingdom) over the 1990–2010 period.

The authors find that dependence and tail dependence are asymmetric for most countries in a negative-correlation regime but symmetric for all countries in a positive-correlation regime. Exchange rate exposure effects or portfolio-rebalancing effects dominate for most countries on most occasions.

11.4 Papers for Further Reading

Below are lists, by category and without comment, of examples of additional papers that may be of value to those interested in either additional depth or possible alternative modeling approaches.

11.4.1 Interest Rates and Credit Risk

Fair, R. C., and B. G. Malkiel, "The Determination of Yield Differentials between Debt Instruments of the Same Maturity," *Journal of Money, Credit, and Banking* 3:4, 1971, pp. 733–749.

Fama, E. F., "The Information in the Term Structure," *Journal of Financial Economics* 13, 1984, pp. 509–528.

Ho, T. Y., and S.-B. Lee, "Term Structure Movements and Pricing Interest Rate Contingent Claims," *Journal of Finance* 41, 1986, pp. 1011–1029.

Brennan, M. J., and E. Schwartz, "A Continuous Time Approach to the Pricing of Bonds," *Journal of Banking and Finance* 3, 1979, pp. 133-155.

Fama, E. F., and R. R. Bliss, "The Information in Long-Maturity Forward Rates," *American Economic Review* 77, 1987, pp. 680–692.

Fama, E. F., "Term Structure Forecasts of Interest Rates, Inflation, and Real Returns," *Journal of Monetary Economics* 25, 1990, pp. 59–76.

Hull, J., and A. White, "Pricing Interest-Rate-Derivative Securities," *Review of Financial Studies* 3:4, 1990, pp. 573–592.

Pennacchi, G. G., "Identifying the Dynamics of Real Interest Rates and Inflation: Evidence Using Survey Data," *Review of Financial Studies* 4, 1991, pp. 53–86.

Hull, J., and A. White, "Numerical Procedures for Implementing Term Structure Models II: Two-Factor Models," *Journal of Derivatives* 2:2, 1994, pp. 37–48.

Black, F., "Interest Rates as Options," *Journal of Finance* 50, 1995, pp. 1371–1376.

Black, F., and P. Karansinski, "Bond and Option Pricing when Short Rates are Lognormal," *Financial Analysts Journal* 47, 1991, pp. 52-59.

Jarrow, R. A., D. Lando, and S. M. Turnbull, "A Markov Model for the Term Structure of Credit Risk Spreads," *Review of Financial Studies* 10:2, 1997, pp. 481–523.

Risa, S., "Nominal and Inflation Indexed Yields: Separating Expected Inflation and Inflation Risk Premia," Columbia University Working Paper, 2001.

Yan, H., "Dynamic Models of the Term Structure," *Financial Analysts Journal*, 57:4, 2001, pp. 60–76.

Chapman, D. A., and N. D. Pearson, "Recent Advances in Estimating Term-Structure Models," *Financial Analysts Journal* 57:4, 2001, pp. 77–-95.

Takamizawa, H., and I. Shoji, "Modeling the Term Structure of Interest Rates with General Short-Rate Models," *Finance and Stochastics* 7:3, 2003, pp. 323–335.

Takamizawa, H., and I. Shoji, "On the Accuracy of the Local Linear Approximation for the Term Structure of Interest Rates," *Quantitative Finance* 4:2, 2004, pp. 151–157.

Gorovoi, V., and V. Linetsky, "Black's Model of Interest Rates as Options, Eigenfunction Expansions, and Japanese Interest Rates," *Mathematical Finance* 14, 2004, pp. 49–78.

Diebold, F. X., and C. Li, "Forecasting the Term Structure of Government Bond Yields," *Journal of Econometrics* 130, 2006, pp. 337–364.

Parnes, D., "Time Series Patterns in Credit Ratings," *Finance Research Letters* 4:4, 2007, pp. 217–226.

Kim, D. H., and K. J. Singleton, "Term Structure Models and the Zero Bound: An Empirical Investigation of Japanese Yield," *Journal of Econometrics* 170, 2012, pp. 32–49.

Gouriéroux, C., and A. Monfort, "Linear-Price Term Structure Models," *Journal of Empirical Finance* 24, 2013, pp. 24–41.

Xiang, J., and X. Zhu, "A Regime-Switching Nelson-Siegel Term Structure Model and Interest Rate Forecasts," *Journal of Financial Econometrics* 11:5, 2013, pp. 522–555.

Dai, Q., and K. J. Singleton, "Specification Analysis of Affine Term Structure Model," *Journal of Finance* 55:5, 2000, pp. 1943–1978.

Arvanitis, A., J. Gregory, and J.–P. Laurent, "Building Models for Credit Spreads," *Journal of Derivatives* 6:3, 1999, pp. 27–43.

Kijima, M., and K. Komoribayashi, "A Markov Chain Model for Valuing Credit risk Derivatives," *Journal of Derivatives* 6:1, 1998, pp. 97–108.

Lando, D., "On Cox Processes and Credit Risky Securities," *Review of Derivatives Research* 2, 1998, pp. 99–120.

Lando, D., *Credit Risk Modeling,* Princeton, NJ: Princeton University Press, 2004.

Cont, R., and P. Tankov, *Financial Modelling with Jump Processes,* Boca Raton, FL: Chapman & Hall/CRC, 2004.

Schonbucher, P. J., *Credit Derivatives Pricing Models,* Chichester, UK: Wiley, 2003.

Christensen, J. H. E., F. X. Diebold, and G. D. Rudebusch, "The Affine Arbitrage-Free Class of Nelson–Siegel Term Structure Models," Federal Reserve Bank of San Francisco Working Paper 2007–20, 2010.

Christensen, J. H. E., and G. D. Rudebusch, "Estimating Shadow-Rate Term Structure Models with Near-Zero Yields," Federal Reserve Bank of San Francisco Working Paper 2013–07, 2013.

Diebold, F. X., and G. D. Rudebusch, *Yield Curve Modeling and Forecasting: The Dynamic Nelson–Siegel Approach*, Princeton, NJ: Princeton University Press, 2013.

Brigo, D., and F. Mercurio, *Interest Rate Models—Theory and Practice: With Smile, Inflation, and Credit,* 2nd ed., Berlin: Springer, 2007.

Rebonato, R., *Modern Pricing of Interest-Rate Derivatives,* Princeton, NJ: Princeton University Press, 2002.

Rebonato, R., K. McKay, and R. White, *The SABR/LIBOR Market Model,* Chichester, UK: Wiley, 2009.

Duffie, D., and K. Singleton, *Credit Risk,* Princeton, NJ: Princeton University Press, 2003.

Duffie, D., and R. Kan, "A Yield-Factor Model of Interest Rates," *Mathematical Finance* 6, 1996, pp. 379–406.

Jamshidian, F., "An Exact Bond Option Formula," *Journal of Finance* 44:1, 1989, pp. 205–209.

Corb, H., *Interest Rate Swaps and Other Derivatives,* New York: Columbia University Press, 2012.

Filipovic, D., *Term Structure Models: A Graduate Course,* Dordrecht, Netherlands: Springer Finance, 2009.

Piazzesi, M., "Affine Term Structure Models," in *Handbook of Financial Econometrics,* edited by Y. Aït-Sahalia and L. Hansen, Amsterdam: North-Holland, 2005.

Ang, A., and M. Piazzesi, 2003, "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables," *Journal of Monetary Economics* 50, 2003, pp. 745–787.

Merton, R. C., "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance* 29:2, 1974, pp. 449–470.

Calvet, L. E., A. J. Fisher, and L. Wu, "Staying on Top of the Curve: A Cascade Model of Term Structure Dynamics," *Journal of Financial and Quantitative Analysis* 53:2, 2018, pp. 937–963.

11.4.2 Equity Pricing and Modeling

Black, F., and M. Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81:3, 1973, pp. 637–654.

Boyle, P. P., "Options: A Monte Carlo Approach," *Journal of Financial Economics* 4:3, 1977, pp. 323–338.

Bru, M.-F., 1991, "Wishart Processes," *Journal of Theoretical Probability* 4:4, 1991, pp. 725–751.

Bollerslev, T., V. Todorov, and S. Z. Li, "Jump Tails, Extreme Dependencies, and the Distribution of Stock Returns," *Journal of Econometrics* 172, 2013, pp. 307–324.

Bollerslev, T., "Generalized Autoregressive Conditional Heteroscedasticity," *Journal of Econometrics* 31:3, 1986, pp. 307–327.

Duffie, D., J. Pan, and K. Singleton, "Transform Analysis and Asset Pricing for Affine Jump-Diffusions," *Econometrica* 68:6, 2000, pp. 1343–1376.

Merton, R. C., "Option Pricing When Underlying Stock Returns Are Discontinuous," *Journal of Financial Economics* 3, 1976, pp. 125–144.

Carr, P., and D. B. Madan, "Option Valuation Using Fast Fourier Transform," *Journal of Computational Finance* 2, 1998, pp. 61–73.

Gatheral, J., *The Volatility Surface: A Practitioner's Guide,* Hoboken, NJ: Wiley Finance, 2005.

Kahl, C., and P. Jackel, "No-So-Complex Logarithms in the Heston Model," *Wilmott Magazine* 19, 2005, pp. 94–103.

Lee, R., "Option Pricing by Transform Methods: Extension, Unification, and ErrorControl," *Journal of Computational Finance* 7:3, 2004, pp. 51–86.

Lewis, A. L., *Option Valuation under Stochastic Volatility, with Mathematica Code,* Newport Beach, CA: Finance Press, 2000.

Albrecher, H., P. Mayer, W. Schoutens, and J. Tistaert, "The Little Heston Trap," *Wilmott Magazine,* January 2007, 83–92.

Andersen, L. B. G., "Simple and Efficient Simulation of the Heston Stochastic Volatility Model," *Journal of Computational Finance* 11:3, 2008, pp. 1–42.

Black, F., "The Pricing of Commodity Contracts," *Journal of Financial Economics* 3:1–2, 1976, 167–179.

Broadie, M., and Ö. Kaya, 2006, "Exact Simulation of Stochastic Volatility and Other Affine Jump Diffusion Processes," *Operations Research* 54:2, 2006, pp. 217–231.

Da Fonseca, J., M. Grasselli, and C. Tebaldi, 2008, "A Multifactor Volatility Heston Model," *Quantitative Finance* 8:6, 2008, pp. 591–604.

Gouriéroux, C., and R. Sufana, "Derivative Pricing with Wishart Multivariate Stochastic Volatility," *Journal of Business & Economic Statistics* 28:3, 2010, pp. 438–451.

11.4.3 Other Literature

Redington, F. M., "Review of the Principles of Life-Office Valuations," *Journal of the Institute of Actuaries* 78:3, 1952, pp. 286–340.

Wilkie, A. D., "A Stochastic Investment Model for Actuarial Use," *Transactions of the Faculty of Actuaries* 39, 1984, pp. 341–403.

Hartmann, P., S. Straetmans, and C. G. de Vries, "Asset Market Linkages in Crisis Periods," *Review of Economics and Statistics* 86:1, 2004, pp. 313–326.

Venter, G., J. Barnett, R. Kreps, and J. Major, "Multivariate Copulas for Financial Modeling," *Variance* 1:1, 2007, pp. 103–119.

Wang, Y.-C., J.-L. Wu, and Y.-H. Lai, "A Revisit to the Dependence Structure between the Stock and Foreign Exchange Markets: A Dependence-Switching Copula Approach," *Journal of Banking and Finance* 37:5, 2013, pp. 1706–1719.

Lipton, A., 2001, *Mathematical Methods for Foreign Exchange: A Financial Engineer's Approach,* Singapore: World Scientific, 2001.

Lipton, A., "The Volatility Smile Problem," *Risk Magazine* 15, 2002, pp. 61–65.

Wystup, U., *FX Options and Structured Products,* 2nd ed., Chichester, UK: Wiley, 2006.

Bates, D. S., "Maximum Likelihood Estimation of Latent Affine Processes," *Review of Financial Studies,* 19:3, 2006, pp. 909–965.

Sims, C. A., "Macroeconomics and Reality," *Econometrica* 48:1, 1980, pp. 1–48.

Kerkhof, J., "Inflation Derivatives Explained: Markets, Products, and Pricing," Fixed Income Quantitative Research, Lehman Brothers, 2005.

Hamilton, J. D., *Time Series Analysis,* Princeton, NJ: Princeton University Press, 1994.

Shreve, S. E., *Stochastic Calculus for Finance II, Continuous-Time Models,* New York: Springer, 2000.

Alfonsi, A., *Affine Diffusions and Related Processes: Simulation, Theory, and Applications,* Cham, Switzerland: Springer International, 2015.

Alfonsi, A., 2010, "High Order Discretization Schemes for the CIR Process: Application to Affine Term Structure and Heston Models," *Mathematics of Computation* 79:269, 2010, pp. 209–237.

Engle, R., "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation, *Econometrica* 50:4, 1982, pp. 987–1007.

11.5 Suggestions on Where to Look for Future Research

When it comes to staying abreast of research developments in the fields that are relevant to ESGs, and specifically to property/casualty applications, the primary sources are the research databases published on the Casualty Actuarial Society website. In particular, the CAS's Database of Actuarial Research Enquiry (DARE) allows a user to search—via keywords and/or dates and/or publications—for relevant citations (and, where available, links) to papers.

An alternative to searching is to employ the CAS taxonomy, which allows one to browse the literature by subject matter.

Whatever approach to DARE is used, the database itself is full of citations of relevance to actuarial work and across numerous publications. The value of DARE is predicated on the dedication and thoroughness of CAS volunteers to identify papers, appearing in any publication, that may have relevance to other CAS members. For ESG purposes, that often means papers in specialized economics and finance journals. Because of this, we encourage the CAS to continue making updates to this database a priority and to make volunteers' ability to identify and report a relevant paper as simple and straightforward as possible.

One can also do occasional Internet searches and even periodically check the tables of contents of certain journals for candidate papers. Even if the full text of articles is behind a firewall, journals will generally make the titles and abstracts of papers available when (or sometimes even before) they are formally published. Here are a few of the journals in economics and finance that may be worth monitoring:

- *American Economic Review*
- *Quarterly Journal of Economics*
- *Journal of Finance*
- *Journal of Financial Economics*
- *Financial Analysts Journal*