

The Importance of Being Transparent: Theory and Implementation Specifics for an Integrated Risk Measurement Model for P&C Companies

Tim Freestone
Roman Gutkovich
Alex Tartakovski

Presented at
Enterprise Risk Management Symposium
Society of Actuaries

Chicago, IL

April 23–26, 2006

Copyright 2006 by the Society of Actuaries.

All rights reserved by the Society of Actuaries. Permission is granted to make brief excerpts for a published review. Permission is also granted to make limited numbers of copies of items in this monograph for personal, internal, classroom or other instructional use, on condition that the foregoing copyright notice is used so as to give reasonable notice of the Society's copyright. This consent for free limited copying without prior consent of the Society does not extend to making copies for general distribution, for advertising or promotional purposes, for inclusion in new collective works or for resale.

Abstract

While most property and casualty (P&C) insurance companies have preferred to implement risk management strategies centered on accounting statements and employing ad-hoc rules and ratios, the financial industry successfully embraced value-at-risk (VaR) methodology based on cash flow analysis (RiskMetrics®, CreditMetrics®, etc.). The historical failure of the insurance industry to adopt a modern risk framework has led companies to rely on such capital adequacy standards as risk-based capital (RBC)—a methodology that offers almost no assistance to management regarding such paramount issues as prudent levels of operational capital, sources of risk, appropriate rates of financial return, modeling of future possible economic scenarios and survival under various market stress conditions.

We present a VaR-based platform that allows for integration of assets and liabilities in a rigorous risk management framework. This Enterprise Risk Model (ERM) measures all the major risks faced by P&C companies: insurance risk, interest rate risk, equity risk, credit risk, foreign exchange and operating risk.

This presentation covers in detail ERM's characteristics:

- ERM is closely based on the VaR methodology of RiskMetrics® and CreditMetrics®. ERM extends this methodology to the long-term risks common to the insurance industry. While RiskMetrics' framework is focused on short-term (3–10 days) trading VaR, we employ a much longer horizon of one year.
- ERM methodology incorporates the correlation structure of assets and liabilities to provide enterprise-wide, fully integrated risk analysis.
- ERM risk factors are estimated from the latest market data. In order to ensure stability and statistical significance of the estimates, ERM applies various calibration techniques, such as principal component analysis (PCA) and random matrix theory.
- ERM employs a combination of Monte-Carlo and Quasi Monte-Carlo simulations. The use of the Quasi Monte-Carlo technique guarantees speed and high accuracy of the simulation output.
- ERM estimates the risk of the whole enterprise and of each business segment, domestic and international.

- ERM calculates various risk measures of the net worth uncertainty, including VaR, Incremental VaR (IVaR), Expected Shortfall (TVaR), standard deviation and downside risk. The rate of convergence is significantly accelerated with the help of importance sampling and robust L-estimates.
- ERM allocates capital by apportioning the risk and diversification benefits to each business segment according to its IVaR share of the total VaR.
- ERM calculates risk-adjusted return on capital (RAROC) and economic capital by business segment.
- ERM provides stress testing, such as a stock market crash, and scenario testing, such as M&A, divestiture, yield curve steepening, or inflation pick-up.

Finally, we apply this advanced risk technology to several companies to reveal risk and capital issues that cannot be identified with conventional industry risk technologies. The authors will demonstrate how this new technology can be used to assist companies in maximizing capital efficiency and migrating from a risk management environment that is governed almost exclusively by ad-hoc rules of thumb to one that is governed by true risk management principles.

About the Authors

Tim Freestone is a managing director at Seabury Analytic LLC. Prior to Seabury, Mr. Freestone was a managing director at KPMG Peat Marwick and was responsible for directing that firm's financial advisory and analysis, valuation and M&A advisory for the insurance industry. Prior to KPMG, Mr. Freestone served as a partner at Mercer Management Consulting where he directed Shareholder Value Analysis Advisory across all industries.

Mr. Freestone has more than 15 years of experience at a senior consulting and advisory level dealing in the areas of shareholder value, corporate valuation, strategic analysis and M&A advisory. Prior to his consulting activities, Mr. Freestone was the director of capital markets for a leading insurance company and the president and CEO of a Prudential Insurance Financial subsidiary.

Mr. Freestone holds a Bachelor of Science from the University of California, at Berkeley.

Dr. Roman Gutkovich is a vice president at Seabury Analytic. Dr. Gutkovich has extensive business and technology experience in both software development and market risk management. He has expertise in the development of risk management software for the leading financial institutions.

Prior to joining Seabury, Dr. Gutkovich worked as a senior technical officer in JPMorganChase, where he was responsible for the design, implementation, and support of an enterprise VaR system. The system performed a firm-wide risk valuation for variety of fixed income, derivative, and equity products. He was involved in the full project life cycle: from collecting user's requirement, to writing specifications, to developing and implementing valuation methodologies.

Dr. Gutkovich has also worked at UBS, where he developed analytical tools for pricing fixed income instruments such as interest rate and FX swaps used by traders for daily trading activities.

Dr. Gutkovich has worked as a senior software developer at BARRA Inc., where he was responsible for the design and implementation of equity portfolio risk management system. He developed analytical software based on a BARRA proprietary multi-factor risk model.

He holds a Master's Degree in applied math and a Ph.D. in fluid mechanics from the St. Petersburg Marine Technical University, as well as Master's Degree in computational finance from Carnegie Mellon University.

Dr. Alex Tartakovski is a vice president at Seabury Analytic. Dr. Tartakovski's expertise is in quantitative risk analysis and modeling for the financial and insurance industries. He has extensive experience in modeling and pricing financial, commodity and insurance derivatives, credit risk analysis and the development of sophisticated, VaR-based portfolio risk management systems.

Prior to joining the Seabury Group in September 2003, Dr. Tartakovski was vice president of research for XL Weather and Energy. He was responsible for modeling activities focused on a wide range of structured derivatives for power, gas, and weather markets. He developed and implemented unique integrated pricing and portfolio risk management systems for multiple-trigger power derivatives and weather deals in support of XL's trading operations.

At Enron Global Markets and ACE USA PowerBacker group, Dr. Tartakovski was heading quantitative research and modeling of power and operation risk derivatives and insurance.

Prior to practicing as a quantitative financial analyst, Dr. Tartakovski worked as an actuary for CIGNA P&C. He was responsible for pricing and statistical support at the Special Risk Facility.

After earning his Ph.D. in theoretical physics from Moscow State University in 1989, Dr. Tartakovski continued his academic research in physics of stochastic processes and superconductivity at the Argonne National Laboratory (Chicago) and University of Florida until 1996. He has published numerous papers in international physics journals and has delivered presentations at physics conferences and research institutions.

1. Introduction

For most of the last 30 years, risk management has been compartmentalized in narrowly defined areas. Market risk was focused on traded instruments and confined to a very short-term horizon. Credit risk was split into two disconnected sub-disciplines of originating decisions and on-going assessments—neither integrated with other aspects of risk.

The asset and liability management (ALM) risk in major banks and insurance companies was confined to short-term forecasts for financial accounting-based net income and cash flow. Using a small number of scenarios, net income was tested for sensitivity to specific changes in the shape and level of yield curves. In commercial banks, at least, both sides of the balance sheet were included in this exercise. In insurance companies, the analysis of income and cash flow was restricted to the investment portfolio—the asset side of the institution only. The actuaries were responsible for the liability side and were not part of the income simulation process.

Financial risk is defined, in most general terms, as uncertainty of results. To measure risk quantitatively, one needs to look at the financial results as stochastic variables and arrive at their distributions, such as Profit and Loss distribution. Traditionally, however, actuaries approached risk deterministically, in terms of pricing and reserve adequacy. When viewed from this perspective, the distinction between the reserve, as the best estimate of the liabilities, and the surplus capital, as the risk margin, was blurred. The actuarial process of reserve setting employed deterministic models, the end result being a point estimate. A single number was supposed to provide not just the best estimate (the expected value) but also afford some degree of risk protection. The protection level, in terms of standard deviations or quantiles of the underlying distribution, was not normally indicated.

The performance measurement and capital allocation were yet other compartments with substantial differences in methodology among similar types of institutions.

Due to a number of key developments in financial theory and computer technology, this “silo” approach to risk management is no longer acceptable. The introduction of the stochastic option pricing models (Black and Scholes (1973) and Merton (1973)) and term structure models of the yield curve (Merton (1973) and Vasicek (1977)), as well as the rapid progress in desktop computational performance, led to the rapid adoption and practical availability of Monte Carlo simulation as a risk management tool. By the late 1980s, these developments culminated in the integrated,

value-at-risk (VaR) based methodology championed by the RiskMetrics team of JPMorgan ((1996) and Mina and Xiao, 2001)). The later development of quantitative credit models (Merton (1974), Jarrow (1995 and 1999) and CreditMetrics (Gupton, Finger, and Bhatia 1997)) completed the analytical integration of credit risk and interest rate risk. In words of van Deventer, Imai, and Mesler (2005), "Market risk, credit risk, and interest rate (ALM) risk all use the same mathematics, the same data, they are caused by the same macro-economic factors, and they impact financial institutions in the same way."

The rapid adoption of the integrated risk management by banks and securities firms has been necessitated, to a large degree, by the pressure from the regulators. The new regulatory requirements, such as Basel II, Solvency II, Financial Services Advisory (FSA), and Australian Prudential Regulation Authority (APRA), focus on integrated risk management based on stochastic models (e.g., see Committee of European Insurance and Occupational Pensions Supervisors (2005), Call No. 11; Financial Services Advisory (2005), PRU 1.4; and The Australian Prudential Regulation Authority (2002, 2005)). The recent accelerated development of such regulations in many countries will make it both mandatory and highly urgent for the global insurance industry to adopt asset and liability risk management on a joint basis. This integrated approach has been standard for the rest of the financial industry for a long time. Common macro factors drive the risk on both sides of the balance sheet. Looking at the risk on only one side of the balance sheet makes it impossible for management, the board of directors, and regulators to have an accurate view of total risk. The very recent introduction of the stochastic actuarial models (Barnett and Zehnwirth (2000), England and Verrall (2002)) clearly shows that the same mathematics so successfully used on the asset side, can and should be applied to the liability risk measurement. There is no doubt that the next few years will witness a profound shift within the insurance industry to the modern integrated risk management.

This paper presents a VaR-based platform that allows for integration of assets and liabilities in a rigorous risk management framework. This Enterprise Risk Model (ERM) measures all the major risks faced by P&C companies: insurance risk, interest rate risk, equity risk, credit risk, foreign exchange, and operating risk.

Although the paper discusses a number of advanced topics, we strove to make this paper accessible to as broad an audience as possible. To this end, the authors adopted a multi-layered approach, wherein each section and subsection begins with a general, relatively non-technical introduction, before proceeding to a more detailed technical discussion; yet more advanced topics are considered in the appendices.

The paper is organized as follows. We introduce definitions and objectives pertinent to risk management in Section 2. In Section 3, we proceed with the description of the analytical framework employed in Seabury Enterprise Risk Model. In Section 4, we introduce the strategic applications of ERM with a case study. In Section 5, we make concluding remarks.

The most difficult technical details related to the implementation issues are delegated to the Appendices A–D. The “know-how” of the kind covered in the appendices is seldom discussed in the literature. It should be noted, though, that in our experience, it is these highly technical details that distinguish between a theoretical development and a practical implementation of an advanced model. In the words of Victor Hugo, “La science est obscure—peut-être parce que la vérité est somber.”¹

2. Risk Management: Definition and Objectives

In the last decade, the definition of risk management has changed dramatically for reasons we outlined in the introduction. The modern definition of risk management can be summarized as follows (van Deventer, Imai, and Mesler (2005)):

Risk management is the discipline that clearly shows management the risks and returns of every major strategic decision at both the institutional and transactional level. Moreover, the risk management discipline shows how to change the strategy in order to bring the risk return trade-off into line with the best long and short-term interests of the institution.

This definition includes the overlapping and inseparable sub-disciplines such as:

- Market risk
- Credit risk
- ALM (interest rate risk)
- Liability risk
- Catastrophe risk
- Operational risk
- Performance measurement and capital allocation

The primary focus of this document is to show how Seabury Enterprise Risk Model implements the best practices of modern risk management in a way that is fully integrated and makes no distinction between these sub-disciplines. These sub-

¹ Science is obscure—maybe because truth is dark (*Facts and Belief*).

disciplines are but different views of the same risk. In what follows, we will prove that the mathematics, data and reporting needs, and information technology infrastructure are all shared by the sub-disciplines of risk management.

3. Seabury Enterprise Risk Model

Seabury Enterprise Risk Model (ERM) was created with the express goal of bringing the best practices of the integrated risk management to the insurance industry. Much more than a simple adoption of the methods pioneered by banks and securities firms, Seabury ERM represents a new methodology specifically designed for the business environment and unique risks of the insurance companies. The developers of Seabury ERM had to overcome significant challenges in order to address the long-term nature of insurance industry and incorporate the insurance liability risks into an integrated stochastic modeling framework.

3.1 Overview

The remaining sections of this paper will describe the functional aspects of the Seabury ERM in somewhat technical terms. We feel it would be a good idea for the reader to have a more general idea of how the model works before delving too deeply into these details. This section presents a general overview of ERM.

3.1.1 ERM is a Single Period VaR Model

One of the principal issues to understand about ERM is that it is a VaR-based single period simulation model, although it produces an entire family of VaR measures including Tail VaR (TVaR)—also known as Expected Shortfall (ES), Marginal VaR, and Incremental VaR (IVaR). VaR is defined as the worst loss that a company may experience over a target period (one year) with a given level of confidence (see Appendix D.1 for more details). If a company's VaR is \$100 million at the 95th level of confidence, this means that there is a 5 percent chance of losing more than \$100 million of net worth over this period of time. VaR is always assessed at what is called the horizon, and the horizon period for ERM is one year.

3.1.2 ERM is a Cash Flow Model

ERM is a cash flow based model that marks all financial investments to market. ERM marks insurance liabilities to model, i.e., the value of the liabilities is the expected value of their future payments over the life of the obligation present valued to the horizon. Earnings in ERM are defined as changes in the value of net worth (assets

minus liabilities). For investments, this is not dissimilar to statutory accounting principles where such factors as realized and unrealized gains/losses are either added to or netted from a firm's surplus. The difference, however, is that statutory accounting principles do not run these credits and debits through the income statement. ERM, on the other hand, converts the firm's income statement from accounting values to market-to-market values so that realistic rates of RAROC can be attributed in the current accounting period.

3.1.3 ERM Produces a Multidimensional Picture of Risk and Risk-Adjusted Performance

ERM is functionally all about producing two kinds of values:

1. Risk measures—the company's total risk, the risk contribution of individual business segments (on both stand-alone and allocated basis), and contribution of individual risk categories (insurance, credit, interest rate, equity, and forex risks),
2. Performance measures, with the focus on valuation of cash flows that go toward the measurement of the firm's RAROC.

Contributing to a firm's risk will be the principal risk categories to which non-life insurance companies are exposed: reserve risk (old business), underwriting risk (new business), equity risk, interest rate risk, credit risk, foreign exchange risk, and catastrophe risk. One can envision that the company's VaR will be influenced over the next year by each one of these risk categories. To understand how this works, one has to consider the cash inflows and cash outflows of an insurance enterprise:

Risk Categories specific to Insurance

- **Underwriting risk** is the risk associated with the new business that will be written over the target period (between the evaluation date and horizon). This risk has two sources: the uncertainty of the new premiums that will be collected over the period and the uncertainty of the future loss payments on the new policies. Income that is allocated to this risk category for RAROC purposes is represented by the present value of the collected premium, minus the present value of the expected loss and expense, plus interest on the economic capital allocated to support the underwriting risk (see Sections 3.2.1.5 and 3.4).

- **Reserve risk:** Business that was written in the past that still remains on the books and for which the company has outstanding reserves is the source of reserve risk. Reserve risk is defined as uncertainty of the loss payments associated with the prior accident years. Income allocated to the reserve risk for RAROC purposes is the present value of duration-matched interest on the insurance reserves plus interest on the economic capital that is allocated to support the reserve risk (see Sections 3.2.1.5 and 3.4).

General Risk Categories

- **Interest rate risk:** Cash flows generated by ERM are discounted in accordance with the term structure of interest rates. ERM simulates individual interest rates for a specified set of maturities and captures both parallel and non-parallel shifts in the yield curve. Interest rate risk is incurred when there is a mismatch between the company's assets and liabilities. The mismatched income is subject to the interest rate risk of the specific time bucket in which it occurs. Cash flows are assigned to time buckets (vertices) in accordance with the RiskMetrics methodology that is covered in the next section.
- **Credit risk** is defined as uncertainty of value due to changes in credit quality. Credit risk in ERM is measured using the CreditMetrics approach which is based on the extended Merton model. Merton's insight was to recognize that the equity in a firm can be considered as a call option on the firm's assets. As a result, the value of the firm's debt can be expressed as the amount of liabilities outstanding reduced by the put option on the assets (the strike being the amount of liabilities). In this option theoretic valuation of debt, bonds become riskier if the return on the underlying equity is weakening, if the maturity is long versus short, if the volatility of the equity is higher rather than lower. The put option reduces the value of debt due to the possibility of default. This basic Merton model can be easily extended to include rating changes (see Section 3.2.1.4).
- **Equity risk:** Equity returns are sensitive to systematic and idiosyncratic factors, all of which are captured in ERM. Systematic factors include sector and country returns. ERM covers 10 equity sectors in all developed countries and many developing countries.
- **Foreign exchange risk:** Income from foreign operations is accounted for in the company's home currency. Foreign income is subject to the same risk factors as income produced in the home currency, plus forex risk.

- **Catastrophe risk:** Losses due to natural catastrophes may have a high impact on the company operations. CAT losses are rare events with high uncertainty which contributes to the overall insurance risk. Also, the large reinsurance receivables generated by the ceded CAT losses are usually the largest single contributor to the credit risk of an insurance company. When aggregating the catastrophe risk with that of the regular losses within a simulation model, the challenge is to properly estimate the impact of the rare CAT events on the overall VaR of the company (see Appendix D.3).

The amount of equity capital that a firm must ultimately carry to support all these risks will depend on its VaR. A large VaR will signal an advanced warning to the firm's leadership that it could potentially lose a level of capital that may impair its operations in the eyes of its stakeholders. This risk is contributed to the VaR from the principal risk categories that have been discussed. But this risk is also mitigated by the degree to which these risk categories are interconnected. For most companies, a significant reduction in the required risk capital (between one-quarter and one-half) could be attributed to correlations inherent in the company's assets and liabilities. So it is extremely important that firms be able to compute diversification benefits with as much accuracy as possible. Issues of correlation are discussed throughout most sections of this paper (especially see Section 3.2.3 and 0).

3.1.4 Virtues of Seabury ERM

The principal virtues that we identify in favor of ERM are:

- ERM is an integrated model that measures all risks and their interdependencies within a single framework
- All assets are processed at the Committee on Uniform Securities Identification Procedures (CUSIP) level and consistently marked-to-market
- All liabilities are marked-to-model and valued on the discounted cash flow basis
- Earnings are defined in terms of net worth (assets minus liabilities) rather than in accounting terms

- The capital that is required to support the company’s risk (and individual business segment risk) is identified at a given level of confidence
- Capital is allocated to support each business segment
- The earnings performance of each business segment is computed to reflect risk-adjusted returns on allocated capital rather than accounting-based measures of earning performance that are inherently flawed.

Hopefully, this description will assist the reader by providing a “big picture” of ERM as subsequent sections delve into substantial detail of how its capabilities are actually executed.

3.2 Analytical Methodology

3.2.1 Framework

- **Cash flows.** Seabury ERM employs the well established cash flow methodology of RiskMetrics®. This methodology far exceeds the “cash flow testing” standards specified by Actuarial Standard of Practice No. 7. While RiskMetrics’ framework is focused on short-term (3–10 days) trading VaR, we employ a much longer horizon of one year.
- **Financial risk factors.** Each asset and liability may have its own idiosyncratic risk, yet be affected by macroeconomic factors. Seabury ERM is a multi-factor model that employs such risk factors as interest rates of different maturity, equity sector returns, and foreign exchange rates in order to capture the effect of the macroeconomic environment.
- **Risk categories.** Bottom-up approach allows for the analysis of the various aspects of the company risk. ERM emphasizes the following categories:
 - Credit risk
 - Interest rate risk
 - Insurance risk
 - Equity risk
 - Currency risk
 - Catastrophe risk.

- **Invested assets.** ERM covers all instruments that are held by investment portfolios such as government, municipal and corporate bonds, asset-backed and mortgage-backed securities (ABS/MBS), common and preferred stocks. All positions are handled at a CUSIP level with the most accurate full valuation algorithms applied.
- **Insurance liabilities.** Seabury ERM employs an extensive database of insurance losses of all U.S. insurance companies and is working to acquire insurance loss data for the other markets. Seabury ERM is based and improves upon the “stochastic reserving” models put forward by Barnett and Zehnwrith (2000), England and Verrall (2002) and others. The incremental accident year losses are subject to the trends due to varying exposure (accident-year trend), development (development-year trend), and inflation (payment-year trend). ERM jointly models the development patterns and development-year dependent loss volatility. Catastrophe risk and credit risk embedded in reinsurance receivables are also explicitly modeled within the integrated framework.
- **Simulation.** Seabury ERM utilizes a Quasi Monte Carlo technique (L’Ecuyer, 2004) supplemented by additional regular Monte Carlo randomization. Due to the use of high performance Quasi Monte Carlo methods based on Korobov’s lattice rules (1959), ERM achieves the speed and rate of convergence impossible with regular Monte Carlo methods (0).

3.2.1.1 Cash Flow Methodology

A portfolio of financial instruments may be broken down into a number of future cash flows associated with each position. However, in the VaR calculation, the large number of combinations of possible cash flow dates leads to the impractical task of computing an intractable number of volatilities and correlations. The RiskMetrics methodology (Mina and Xiao, 2001) drastically simplifies the time structure by mapping each cash flow to a pre-specified set of vertices. In ERM, each U.S.-denominated cash flow is mapped to one or more of the vertices shown below:

$$\leq 1\text{yr} \quad 2\text{yrs} \quad 3\text{yrs} \quad 5\text{yrs} \quad 7\text{yrs} \quad 10\text{yrs} \quad 15\text{yrs} \quad 20\text{yrs} \quad \geq 30\text{yrs} \quad (3.2.1)$$


Mapping a cash flow means splitting it between two adjacent vertices in such a way that both the present value of the cash flow and its sensitivity to the zero rates are preserved. As a result of mapping, a portfolio of instruments is transformed into a portfolio of standard cash flows. **Error! Reference source not found.** shows how the

actual cash flow at year six is split into the synthetic cash flows at the five-year and seven-year vertices.

RiskMetrics documentation ((Mina and Xiao, 2001, pp. 43–45) shows that a payment of USD 1 at time t could be mapped into a payment of W_L at time t_L , a payment of W_R at time t_R (t_L and t_R are the two adjacent vertices around t , $t_L < t < t_R$), and a cash position C :

$$\begin{aligned}
 W_L &= \alpha \frac{t}{t_L} e^{-z_t t} e^{z_L t_L} \\
 W_R &= (1-\alpha) \frac{t}{t_R} e^{-z_t t} e^{z_R t_R} \\
 C &= -\frac{(t-t_L)(t_R-t)}{t_R t_L} e^{-z_t t}
 \end{aligned}
 \tag{3.2.2}$$

Figure 1
Cash Flow Mapping



This rule assumes that the zero rate z_t for maturity t is calculated as a linear interpolation of zero rates z_L and z_R at the vertices, i.e.,

$$z_t = \alpha z_L + (1-\alpha) z_R,
 \tag{3.2.3}$$

where $\alpha = (t_R - t) / (t_R - t_L)$.

3.2.1.2 Financial Risk Factors

As discussed above, the Seabury enterprise risk management platform relies on a multi-factor methodology employed by RiskMetrics risk model. RiskMetrics does not look at the company portfolio as a set of assets which ought to be processed independently, but analyzes it in terms of common risk factors affecting the value of assets. Seabury ERM utilizes such risk factors as interest rates of different maturities for the U.S. dollar (USD) and other currencies, equity sector returns for the United States and other countries, and foreign exchange rates. This multi-factor approach is widely accepted in the financial industry due to its practicality. One of the benefits of this approach is that the assessment of risk that is incorporated into a complex portfolio structure could be reduced to the analysis of a limited number of risk factors; correlations between different asset classes and risk categories will be derived straightforwardly through the exposures to the specific risk factors. Another benefit is that, through the use of a comprehensive set of risk factors, the analyst may model various market environments and evaluate the impact to the company arising from the change in market conditions.

All factors employed by ERM can be observed directly in the market, therefore important factor characteristics such as volatilities and correlations of the returns can be obtained directly from the historical price series through the methods of statistical analysis. The distribution of past returns can then be modeled to provide a reasonable forecast of future factor returns over the required horizon. For each individual instrument, Seabury ERM identifies the set of the specific risk factors which drive the change in the instrument price as well as the exposure to each factor. By generating future scenarios for each risk factor, ERM infers changes in the instrument value and re-prices the total company portfolio accordingly. Such a bottom-up approach possesses a great degree of flexibility and simplifies the broad analysis of the company.

We follow the methodology of the RiskMetrics deriving distributions of parameters for risk factors from the historical logarithmic returns series:

$$r_{t,H} = \ln(P_{t+H}/P_t), \quad (3.2.4)$$

where $r_{t,H}$ denotes the return from time t over the horizon period H to $t + H$ and P is a generalized price which, depending on the risk factor, may represent the Treasury bond price, industry index value, or exchange rate. RiskMetrics advocates the use of the exponentially weighted returns for the estimation of volatility; this schema assigns more weight to the most recent data and limits the effective number of historical returns. While this approach seems appropriate for a short-term horizon, it is not

suitable for a one-year horizon which is considered necessary due to the nature of insurance liabilities. Taking these issues into account, the authors selected the equally weighted volatility estimate, which uses historical data series going back 10 years.

The model for the distribution of future returns is based on the notion that logarithmic returns of risk factors are jointly normally distributed. J. Mina and J. Yi Xiao (2001) outline the arguments that justify the practical use of normal distributions for the problem at hand—fast and accurate estimation of various risk statistics for a portfolio driven by a large number of risk factors.

A practical justification for the normal distribution is the simplicity of its calibration. The univariate normal distribution can be described by two parameters that are easy to calibrate: the mean and standard deviation. Every distributional model has to consider the dependence structure of the returns as well as their stand-alone characteristics. The most important practical advantage of the multivariate normal distribution is that its dependence structure is uniquely defined by a correlation matrix.

Normal distributions, though, cannot adequately describe rare events which result in big losses, such as a catastrophe loss. Modeling such an event would require a skewed distribution with a heavy negative tail. This issue will be discussed in more detail in the CAT risk section.

It is essential that all generalized prices used for the returns and subsequent correlation estimates must be denominated in USD. Use of a single currency across all the instruments either domestic or foreign in order to produce a correlation structure simplifies a transition from one reporting currency to another (rebasings) with no recalculations of the correlation matrix and principal components required.

To illustrate this concept, let us consider a return on 10-year Treasury strip. The price of the instrument is denoted as P^{10y} . As long as the base currency is USD, the return is a relative change in the domestic price:

$$r_{t,H}^{10y\ USD} = \ln\left(\frac{P_{t+H}^{10y\ USD}}{P_t^{10y\ USD}}\right) \quad (3.2.5)$$

If the company operates in Europe and reports its earnings and risks in EUR, then from its perspective, a holding of the U.S. Treasury instrument in the investment portfolio would be subject to EUR/USD exchange rate risk. Even if domestic price of the Treasury stays the same over the horizon period, a change in EUR/USD rate can increase or reduce its value for a foreign company:

$$\begin{aligned}
r_{t,H}^{10y\text{ EUR}} &= \ln\left(\frac{P_{t+H}^{10y\text{ EUR}}}{P_t^{10y\text{ EUR}}}\right) \\
&= \ln\left(\frac{P_{t+H}^{10y\text{ USD}} \cdot [USD / EUR]_{t+H}}{P_t^{10y\text{ USD}} \cdot [USD / EUR]_t}\right) \\
&= \ln\left(\frac{P_{t+H}^{10y\text{ USD}}}{P_t^{10y\text{ USD}}}\right) + \ln\left(\frac{[USD / EUR]_{t+H}}{[USD / EUR]_t}\right) \\
&= r_{t,H}^{10y\text{ USD}} - r_{t,H}^{[EUR/USD]}
\end{aligned}
\tag{3.2.6}$$

In other words, the Euro-denominated return on the Treasury could be derived from the USD-denominated return by simply subtracting the return on the EUR/USD exchange rate. Since the same is true for any USD-denominated instrument, the rebasing schema could be depicted as follows:

1. Simulate USD-denominated returns for all risk factors
2. Simulate CCY/USD returns for a new reporting currency CCY for each scenario
3. Recalculate factor returns by applying simulated foreign exchange rates
4. Apply new factor returns and produce a new company's value for each scenario. Calculate new CCY-denominated risk statistics from the new distribution of the company values.

3.2.1.3 Risk Categories

The bottom-up approach allows for the analysis of the various aspects of a company's risk. Breaking risk down into its subcategories proved very useful for understanding the uncertainties faced by the company and protecting it from potential losses. These subcategories are nothing but different views of the same risk. ERM emphasizes the following categories:

Insurance risk—the uncertainty associated with future payments of insurance liabilities. The factors driving this risk are the size and structure of the insurance business. ERM measures this risk from the analysis of the company's loss triangles and historical underwriting results.

Credit risk—the uncertainty associated with changes in obligor credit quality. On the investment side, this category indicates the potential loss in the net worth the company may experience from the deterioration in the credit quality of its investments assets. On the insurance side, the main component of the credit risk

is the credibility of the reinsurers who may fail to pay on their obligations. ERM measures credit risk from the historical rating upgrade and downgrade records.

Interest rate risk—the uncertainty associated with a change in interest rates. It measures a change in the net value of fixed income instruments and insurance liabilities resulting from the potential fluctuations in interest rates. In most cases, variations in future interest rates impact present value through the adjustment of discounting factors. But for some instruments like callable bonds or ABS/MBS, changing interest rates may impact the projected cash flows. ERM estimates interest rate risk parameters from the historical variations in government yield curves.

Equity risk—the uncertainty associated with the stock market volatility. ERM applies the historical experience of stock market movements to the investment portfolio in order to assess a potential loss in equity positions.

Currency risk—the uncertainty associated with fluctuations in exchange rates. Measures potential loss for the company which is doing business abroad and/or keeps instruments denominated in foreign currency in its investment portfolio. ERM estimates this risk from the historical variations in foreign exchange rates.

Catastrophe risk—the uncertainty associated with the impact that natural catastrophes may have on a value of the company. We will discuss this risk in the Catastrophe Risk subsection of Section 3.2.1.5.

It is important to note that risk subcategories are not independent. Even though each particular category represents a distinct aspect of the enterprise risk, they are driven by the common set of factors which, in turn, are closely correlated to each other. Exchange rates are not independent from the interest rates of the participating currencies, and the credit rating of a company may be closely related to its stock performance. As a result, the total risk of the company may be significantly lower than the sum of the individual risks. The difference between the two indicates the level of correlation that exists between risk categories and is usually referred to as the diversification benefit. Unlike many risk management systems that rely on a top-down approach, ERM does not make any assumptions about correlations between risk categories and the resulting risk reduction. Estimates of correlation arise logically from the bottom-up analyses. We will discuss this issue in more detail in the context of risk aggregation capabilities offered by ERM.

3.2.1.4 Investment Portfolio

Common stock

Seabury ERM employs the linear regression model assuming that the standardized log return of the firm's value, r^e , is the weighted average of two standardized returns, namely, the industry return, r_I , and the firm-specific return, ε :

$$r^e = w_I r_I + \sigma \cdot \sqrt{1 - w_I^2} \varepsilon \quad (3.2.7)$$

where $\varepsilon \sim N(0,1)$ and volatility σ could be derived from the historical stock prices

The practical interpretation of the above equation is that the firm's return can be sufficiently explained by the index return of the industry classification to which the firm belongs, with a residual part that can be explained solely by information unique and specific to the firm. Firm-specific risk can generally be considered to be a function of company asset size. Larger companies tend to have smaller firm-specific risk while smaller companies, on the other hand, tend to have larger firm-specific risk. According to JPMorgan's CreditManager, the firm-specific risk follows the logistic curve:

$$FirmSpecificRisk = \frac{1}{2(1 + Assets^{0.4884} \times e^{-12.4739})} \quad (3.2.8)$$

with *Assets* being the total assets in USD. For asset size of \$1 billion, firm-specific risk is 0.46, implying $w_I = 0.54$. For asset size of \$100 billion, $w_I = 0.75$. Each simulation scenario produces a realization for all index returns and specific returns for all stocks positions, thus assigning new value for the equity portfolio.

Risk-free bonds

ERM views a risk-free coupon-paying bond as a deterministic stream of future cash flows. Applying cash-flow mapping procedure as described in Section 3.2.1.1 above, ERM maps the future payments into individual vertices denoted as W_t . To calculate horizon value of the bond, a cash flow at every vertex is discounted using the appropriate domestic risk-free curve. The authors selected to use U.S. and foreign synthetic zero curves provided by Bloomberg® for discounting cash flows. This procedure can be repeated for all coupon-paying bonds held in the company's portfolio. Cash flows from individual instruments are aggregated into the suitable maturity vertices.

The market value of the bond portfolio becomes

$$V_h = \sum_{t_i \in \text{vertices}} \left(\sum_j W_{t_i}^j \right) e^{-z_{t_i} t_i} + \sum_j C^j, \quad (3.2.9)$$

where index j denotes the individual bond; z_{t_i} is the zero rate with the maturity t_i ; and C is a cash position produced from the mapping algorithm. For each scenario, ERM generates the array of simulated zero rates, substitutes them into the equation, and calculates a new portfolio value. Simulations therefore result in a distribution of the projected portfolio values.

For cash flows that are within the time horizon, ERM takes the conservative approach and assumes that the cash flow earns the interest at the constant risk-free short-term rate and so the present value at the horizon is just the sum of the accrued cash flows.

Risky bonds

Unlike risk-free bonds, where future cash flows are deterministic and the projected market value is only subject to interest rate uncertainty, the risky bonds have exposure to default risk as well. To capture this risk, ERM models the change in the credit quality of the bond over the specified horizon through the use of a transition probability matrix—rules that show how credit ratings migrate over unit time intervals. Whether the credit rating of the bond improves, deteriorates, or stays the same, the market value of the instrument adjusts accordingly. TABLE 1 below shows transition probabilities and resulting values of a hypothetical BBB bond over a one-year period.

TABLE 1
Year-End Values after Credit Rating Migration from BBB

Current Rating	Possible Future Rating	Probability	Resulting Value
	AAA	0.02%	\$101.69
	AA	0.33%	\$101.47
	A	5.95%	\$101.03
BBB	BBB	86.9%	\$100.00
	BB	5.30%	\$94.86
	B	1.17%	\$91.21
	C	0.12%	\$77.77
	D	0.18%	\$47.54

ERM employs the CreditMetrics' asset value model, which links the return on a company's stock with its probability of being upgraded or downgraded within the examined period of time. The asset value model assumes that the one-year return is normally distributed and that the bond's rating changes to a new value when the normalized return drops below or jumps above the respective threshold as illustrated by the following chart. The thresholds could be calculated from transition probabilities as dictated by a normal distribution:


$$\begin{aligned} \text{Prob}(\text{Default}) &= \text{Prob}(r < Z_D) = \Phi(Z_D) \\ \text{Prob}(\text{CCC}) &= \text{Prob}(Z_{\text{CCC}} < r < Z_D) = \Phi(Z_C) - \Phi(Z_D), \end{aligned} \quad (3.2.10)$$

which yields for the threshold S :

$$Z_S = \Phi^{-1} \left[\sum_{l=D}^S \text{Prob}(l) \right], \quad S = \text{CCC}, B, \dots, \text{AAA}. \quad (3.2.11)$$

ERM thus evaluates the credit risk embedded in a corporate bond by simulating a return on the issuing firm's stock price. For a portfolio of risky bonds, the co-movements in credit migrations of different bonds are captured through the simulation of correlated returns of the corresponding stock prices.

Figure 2
Distribution of Asset Returns with Rating Change Thresholds



Interest rate risk assessment for a risky bond is analogous to that for a risk-free bond explained in the previous section.

Many sovereign and corporate bonds carry a call provision which grants the issuer an option to retire (or “call”) the bond prior to its maturity. The callable bond value therefore equals the “optionless” bond value, less the call option value. Since the option value depends primarily on the current interest rates, and changes along with the changes in a yield curve, this value should be recalculated for each scenario whenever interest rates fluctuate. ERM utilizes the Hull-White model of interest rate evolution to calculate a value of the callable bond for each simulation scenario.

ABS/MBS

When calculating risk for ABS/MBS securities, one must take into consideration that these instruments (unlike bonds) carry a prepayment provision granted to the borrower. This means ABS/MBS may be fully or partially prepaid by the borrower at any time he or she selects. This option has a significant importance; its valuation becomes the integral part of ABS/MBS full valuation. The proper option valuation requires the use of a prepayment model, which utilizes a wide range of historical data and analyzes different economic factors. Implementing such a model would go beyond the scope of the ERM platform, instead, ERM relies on key rate durations and the expected horizon prices provided by the user or calculated with specialized software like that developed by CMS BondEdge® or Citigroup YieldBook®, or other vendors. Key rate durations (also called option-adjusted durations) reflect the change in instrument value with respect to a small change in an interest rate for a specific maturity bucket (key rate). Part of this value change comes from the variation in prepayment speed, driven mainly by interest rates; therefore this method captures prepayment risk rooted into the total instrument risk. ERM calculations also include convexity, which guarantees a second order of accuracy. Effective duration and convexity are estimated from the parallel shift of the entire yield curve.

For i th scenario change in ABS/MBS market value at the horizon becomes

$$MV_h^i = MV_h^0 \cdot \left(1 - EDUR \cdot \overline{\Delta z}^i + \frac{1}{2} ECNVX \cdot \overline{\Delta z}^i{}^2 - \sum_{j=1}^N KRD_j \cdot (\Delta z_j^i - \overline{\Delta z}^i) \right), \quad (3.2.12)$$

where

$$\overline{\Delta z}^i = \frac{1}{N} \sum_{j=1}^N \Delta z_j^i, \quad (3.2.13)$$

is an averaged parallel shift, Δz_j^i is a fluctuation of a zero key rate of j th maturity bucket around its central horizon value for i th scenario. In Equation (3.2.12), the effective duration $EDUR$ and convexity $ECNVX$ are the first and second degree order changes in the ABS/MBS price with respect to a small parallel shift in the yield curve. Like KRD s,

they encompass the impact that the changing interest rates may have on prepayment speed.

Preferred stock

Usually insurance companies hold only a small portion of the investment portfolios (up to 2 percent) in preferred stocks. More than 90 percent of them are of callable non-convertible cumulative types which are similar to corporate bonds. To assess their risk, ERM treats preferred stock as callable corporate bonds with the identical maturity, face value, coupon percentage and rating. If stock does not have a maturity, ERM assigns it the longest maturity term available for a given currency. Given that preferred stocks are subordinate to even the least senior bonds, ERM sets their recovery rate to zero.

3.2.1.5 Insurance Liabilities

Introduction

Insurance companies assume risks of other entities (individuals and companies) in return for payments of premium. While premium is normally paid up front, before the inception of an insurance contract, loss payments by the insurance company will not take place until after an insured event occurs and is reported. Even when the insured event is reported, the exact magnitude of the loss payments due to the insured (“ultimate loss”) may not be known for quite some time. The determination of this amount (known as “adjustment process”) may be lengthy and involve litigation. As a result, the loss payments attributable to any given policy may occur over a period of time even after the coverage period ends, potentially in the course of many years (e.g., over a lifetime of an insured individual)—a process known as “loss development.” Thus, loss payments become a long-term liability for the insurance company and give rise to uncertain future cash flows.

Insurance companies are required to establish reserves, i.e., money set aside, for the future loss payments. The reserves consist of case reserves, as set by the case adjusters: bulk reserves—these judgmental adjustments to case reserves are established by actuaries on an aggregate basis, and Incurred But Not Reported (IBNR) reserves—the latter being an actuarial allowance for yet unknown events. The sum of already paid losses and reserves for future payments is referred to as incurred losses.

The insurance risk modeling in ERM requires analysis and simulation of cash flows stemming from claim payments. From the accounting point of view, insurance

liabilities are represented by the incurred losses, while the risk management perspective requires analysis of the uncertainty in the actual cash flows which are represented by paid losses. A change in the value of the incurred losses reflects a change in the internal estimate of the unpaid loss. A measure of variability of such estimates would not be representative of the intrinsic fluctuations in the paid losses, and, therefore, would not provide a measure of the risk of the associated cash flows. Also, the current practice of reserve setting is not based on statistically sound analysis of claim data and is subject to varying company policies. As a result, in ERM we concentrate on the development and fluctuations of the paid losses. The incurred loss data may be used as a supplement in order to estimate the development beyond the period covered by the paid loss data on Schedule P (10 years).

Loss Triangles and Link Ratios

Historical loss data are conveniently organized in the form of triangles.

TABLE 2
Cumulative Paid Losses

		Development Year									
		0	1	2	3	4	5	6	7	8	9
Accident Year	1981	5,012	8,269	10,907	11,805	13,539	16,181	18,009	18,608	18,662	18,834
	1982	106	4,285	5,396	10,666	13,782	15,599	15,496	16,169	16,704	
	1983	3,410	8,992	13,873	16,141	18,735	22,214	22,863	23,466		
	1984	5,655	11,555	15,766	21,266	23,425	26,083	27,067			
	1985	1,092	9,565	15,836	22,169	25,955	26,180				
	1986	1,513	6,445	11,702	12,935	15,852					
	1987	557	4,020	10,946	12,314						
	1988	1,351	6,947	13,112							
	1989	3,133	5,395								
	1990	2,063									

The actuarial technique routinely used to analyze loss triangles is the “chain ladder link-ratios” method. In this approach, one finds average ratios of cumulative payments in each pair of adjacent development years. For example, the average ratio of payments in column 1 to same accident year payments in column 0 (averaging could be straight, dollar weighted, “excluding high-low,” etc.) would give the development factor between “ages” 0 and 1. From a statistical point of view, this technique is based on the assumption that a cumulative payment in one development year is a predictor of the incremental payment in the next year, which means, in particular, that incremental payments of the same accident year are not independent. On the contrary, statistical

loss data analysis in the literature (Barnett and Zehnwirth, 2000) and Seabury’s own research show that for most real loss development arrays, the incremental payments are independent random variables, and, therefore, standard development factor (link-ratio) techniques are inappropriate.

This finding has important implications for both reserve setting (estimates of the mean) and risk analysis (estimates of uncertainty). First of all, it can be proven that the “chain ladder link-ratios” method produces upward-biased reserve estimates (Korobov, 1959). The consequence for the reserve uncertainty is rather subtle: under the assumptions behind the “chain ladder” approach, the probability distribution of the incremental payments would necessarily widen as a function of the development “age” (similarly to how the volatility of a common stock scales with time). The later development years would then make a bigger contribution to the reserve risk than their contribution to the reserve itself. On the other hand, under the assumption of statistically independent incremental payments, there is no inherent widening of the distribution with time—although the volatility may still increase late in the development because of the low number of claims remaining open.

Statistical Modeling Framework

It is clear from the above discussion that the natural way to analyze and model paid loss data is on the incremental rather than cumulative basis.

**TABLE 3
Incremental Paid Losses**


		Development Year									
		0	1	2	3	4	5	6	7	8	9
Accident Year	1981	5,012	3,257	2,638	898	1,734	2,642	1,828	599	54	172
	1982	106	4,179	1,111	5,270	3,116	1,817	(103)	673	535	
	1983	3,410	5,582	4,881	2,268	2,594	3,479	649	603		
	1984	5,655	5,900	4,211	5,500	2,159	2,658	984			
	1985	1,092	8,473	6,271	6,333	3,786	225				
	1986	1,513	4,932	5,257	1,233	2,917					
	1987	557	3,463	6,926	1,368						
	1988	1,351	5,596	6,165							
	1989	3,133	2,262								
	1990	2,063									

The incremental loss development array is subject to multiple trends acting in different directions: the natural loss development within each accident year (horizontal

direction in TABLE 3), the change in exposure from one accident year to another (vertical direction), and inflation—since calendar years are represented by diagonals in TABLE 3, this last type of trend acts from one diagonal to another. Seabury ERM introduces a modeling framework for the incremental paid loss data that is able to capture all these trends.

The model parameterizes the trends in each of the three directions—development years, accident years and payment/calendar years. In what follows, the development years are denoted by j , $j = 0, 1, 2, \dots, s-1$; accident years by i , $i = 1, 2, \dots, s$; and payment years by t , $t = 1, 2, \dots, s$.

The payment year variable t can be expressed as $t = i + j$. This relationship implies that both development and accident year trends are projected onto payment year trends.



In its most general form, the model can be written as:

$$P_{i,j} = \exp\left(\alpha_i + \sum_{k=1}^j \gamma_k + \sum_{t=1}^{i+j} \iota_t\right) + \varepsilon_{i,j} \quad (3.2.14)$$

Here, $P_{i,j}$ is the incremental payment amount for accident year i and development year j —the payment takes place in payment year $i + j$; $\varepsilon_{i,j}$ is zero-mean random error (not necessarily normally distributed).

The parameters α_i in the accident year direction determine the level from year to year; often the level (after adjusting for exposures) shows little change over many years, requiring only a few parameters. The parameters γ_k in the development year direction represent the trend from one development year to the next. This trend is often linear (on the log scale) across many of the later development years, often requiring only one

parameter to describe the tail of the data. The parameters ι_t in the payment year direction describe the trend from payment year to payment year. If the original data are inflation adjusted, the payment year parameters represent superimposed (social) inflation, which may be stable for many years. This is determined in the analysis. Consequently, the (optimal) identified model for a particular loss development array is likely to be parsimonious. This allows us to have a clearer picture of what is happening in the incremental loss process.

The distribution of variables $P_{i,j}$ in the model (3.2.14) is determined by the distribution of the random term $\varepsilon_{i,j}$. It is well known that insurance losses exhibit long-tail distributions. Accordingly, the natural candidates would be such distributions as lognormal and fat-tail power-law family. Note that the random term in (3.2.14) is additive, so that the incremental payment might become negative. This is actually advantageous, because the real insurance data, due to the practices of salvage and subrogation, do exhibit this kind of anomaly (as evident from the data in TABLE 3).

Implementation and Calibration

The model defined by Equation (3.2.14) is nonlinear and cannot be reduced to a linear regression. Accordingly, the parameters in (3.2.14) should be obtained through a nonlinear regression analysis (general nonlinear minimization). In the current ERM version, we accomplish a less ambitious goal and approximate Equation (3.2.14) by a linear model on a log-scale:

$$y(i, j) \equiv \ln(P_{i,j}) = \alpha_i + \sum_{k=1}^j \gamma_k + \sum_{t=1}^{i+j} \iota_t + \varepsilon_{i,j}. \quad (3.2.15)$$

Note that Equation (3.2.15) does not allow negative incremental payments; therefore, we are forced to drop negative data points from the historical data.

The original equation (3.2.14) has an inherent advantage missing from (3.2.15): any scheme used for minimization the random terms in (3.2.14) would be dominated by the large dollar amounts (usually, early development years), precisely the ones that contribute most into the risk of the future payments. In the log-scale model (3.2.15), if one were to perform an ordinary least-squares regression (OLS), the parameters could be driven by the large variations in the small dollar amounts of the old development age. Therefore, our approximation of (3.2.14) by (3.2.15) makes it necessary to perform a weighted least-squares regression in (3.2.15), the weights being the payments $P_{i,j}$ themselves.

The random term in (3.2.15) is assumed to be normal, so that the incremental payments follow a lognormal distribution. We do not assume, however, that the error terms $\varepsilon_{i,j}$ come from a single distribution. Instead, we model $\varepsilon_{i,j}$ as $N(0, \sigma(P_{i,j}))$, where the standard deviation $\sigma(P_{i,j})$ becomes a function of the incremental payment $P_{i,j}$. The chosen scheme of weighted least-square regression implies that the variances σ_j^2 are inversely proportional to $P_{i,j}$, the proportionality factor being determined by the regression. We, however, make an additional step and assume that $\sigma(P_{i,j})$ is a general non-increasing function of $P_{i,j}$. This assumption is motivated by our research that shows that the payment volatility (on log-scale) remains practically constant over a wide range of payment magnitude, but once the payments drop below a certain threshold, the volatility begins to increase. Therefore, having done the weighted regression as discussed above, we perform an additional non-parametric fit of the squared residuals with a monotonous (more specifically, non-increasing) function of $P_{i,j}$; this least-square fit is weighted, $P_{i,j}$ being the weights again. The resulting function provides the required estimate of variance $\sigma^2(P_{i,j})$:

$$\begin{aligned}\varepsilon_{i,j}^2 &= \sigma^2(P_{i,j}) + v_{i,j}, \\ \sigma^2(P) &\leq \sigma^2(P') \text{ for } P > P'.\end{aligned}\tag{3.2.16}$$

In Equation (3.2.16), $v_{i,j}$ is an error term; if $\varepsilon_{i,j}^2$ happens to be non-increasing as a function of $P_{i,j}$, then all $v_{i,j}$ would be equal to 0.

Theoretically speaking, Equation (3.2.15) is overparameterized by one parameter and exhibits “perfect” multicollinearity: it remains invariant under the transformation

$$t_i = t_i - I, \quad \gamma_i = \gamma_i + I, \quad \alpha_i = \alpha_i + iI,\tag{3.2.17}$$

where I is arbitrary. As a result, the mean level of the inflation over all payment years cannot be determined (it is included in the development factors γ_i); only the deviations from the mean level can be calculated. This situation would change if we knew the accident year exposures so that we could normalize the incremental payments and set all α_i to be equal to each other. Unfortunately, we are unaware of any publicly available information about the exposures or suitable proxies. Accident year premiums, in particular, cannot serve as good proxies for exposures due to the varying premium rates over a business cycle. In the absence of the company’s exposure data, we get rid of multicollinearity by setting

$$t_1 = 0;\tag{3.2.18}$$

this means that the rest of the γ parameters measure the *difference* between inflation in their respective years and that in year 1.

Even though Equation (3.2.15) with restriction (3.2.18) exhibits no theoretical (“perfect”) multicollinearity, it still has too many fitting parameters to be of practical use in forecasting. In particular, it is reasonable to require that the development pattern captured by γ_i represents a smooth curve. Yet, as long as all γ_i are estimated independently, this condition cannot be guaranteed. We will both achieve a parsimonious model suitable for forecasting and ensure smooth development patterns if we significantly reduce the number of parameters in (3.2.15) by setting some of them equal to each other or to zero. For instance, we might require that, for a given line of business, $\gamma_1 = \gamma_2$, $\gamma_3 = \gamma_4 = 0$, $\gamma_5 = \gamma_6 = \dots = \gamma_{s-1}$, and $t_2 = t_3 = \dots = t_s$. We refer to such specifications as a “model structure”. We determine the proper model structure for each line of business based on the industry-wide experience. The concrete values of the parameters remaining in the model and their standard deviations can then be determined for each company; these values can possibly be credibility weighted with the results of the industry cross-sectional analysis.

Because parameters γ_i and t_i represent trends that will accumulate when we set some of them equal, we need to account for a “model risk” by making the parameters themselves normally distributed random variables. In the example above, in each scenario we would need to simulate two γ parameters, one t parameter, and only afterwards all the random terms $\varepsilon_{i,j}$. The mean values and the standard deviations for this simulation are the output of the regression (3.2.15). Note that parameters α_i are not trends; these parameters should simply be set to their regression estimates rather than simulated.

Note that parameterization is such that in any forecast (simulation) we set the future values of parameters α_i , γ_i , and t_i to be the same as in the last year available from the regression.

In order to account for correlations between the insurance lines and the correlations between assets and liabilities, we regress the normalized error terms $\varepsilon_{i,j}/\sigma(P_{i,j})$ against the assets’ principal components and then perform the principal component analysis on the residuals of that regression. As the end result of this analysis, the random term in (3.2.15) is represented as a linear combination of principal components of both assets, PC^{asset} , and liabilities, $PC^{\text{liability}}$, plus an idiosyncratic random term:

$$\varepsilon_{i,j}/\sigma(P_{i,j}) = \sum_{m=1}^M C_{i+j,m}^{\text{asset}} \times PC_{i+j,m}^{\text{asset}} + \sum_{n=1}^N C_{i+j,n}^{\text{liability}} \times PC_{i+j,n}^{\text{liability}} + \tilde{\varepsilon}_{i,j}, \quad (3.2.19)$$

where $\tilde{\varepsilon}_{i,j}$ are i.i.d. $N(0, \sigma)$, M is the number of principal factors on the asset side and N is the number of factors on the liability side.

Underwriting Risk (New Business)

The uncertainty of the loss payments associated with the prior accident years (old business)—usually referred to as reserve risk—is captured by the model introduced in the previous sections. The same model is used in ERM to describe the loss uncertainty of the new business that will be written between today and the horizon. This future business has one more component that we have not covered yet—the uncertainty of the collected premium. Insurance premium collected per unit of risk (premium rate) is subject to market forces; as a result, the insurance industry has gone through well documented business cycles of “hard” and “soft” markets. Even though the uncertainty of the rate forecast between today and the one-year horizon is usually much less than the rate variations over an entire business cycle, the random element in the future premium cannot be removed.

Historically, the insurance business cycles have not coincided with the economic cycles. Nonetheless, the existence of a relationship between premium rates and the economic environment, in particular, the interest rates, is well known (Taylor, 2003; Doherty and Garven, 1995; Grace and Hotchkiss, 1995; Fung et al., 1998). Since most collected premiums get invested into fixed income instruments, a hike in the interest rates will result in greater investment income associated with the new policies. In a competitive market environment, this will result in additional pressure towards lowering the premium charged per unit of risk. This negative correlation between interest rates and premium rates is more pronounced in casualty lines where claims are settled long after the premiums are collected than for short-tailed property lines. Within Seabury ERM, premium rates are modeled as lognormally distributed variables; the correlations between rates in different lines and between rates and the financial risk factors are estimated based on the industry-wide data.

Catastrophe Risk

Losses due to natural catastrophes (CAT losses) are rare events with very high impact. The CAT loss distributions exhibit highly non-normal (fat-tailed) behavior. The company specific distributions depend on the geography of the insured properties and businesses. As a result, such losses are notoriously difficult to model reliably, and

the distributions are provided by just a few vendors who specialize in CAT risk. When the distribution is available, however, it is relatively easy to include the CAT risk into the overall framework of ERM. Within the simulation model, we need to add catastrophic losses drawn from the CAT distribution to the losses generated in our regular model (3.2.15). This procedure will result in adjusted α_i (and possibly γ_k , if we assume that CAT loss development is different from that of the regular losses) in such catastrophic scenarios. In addition, the CAT scenarios will have significant reinsurance receivables, hence additional credit risk (see below). The real difficulty with catastrophic risk within a Monte Carlo approach stems from the necessity to estimate the risk measures dominated by rare events. In Appendix D, we show how application of importance sampling and robust estimators can overcome these problems.

Reinsurance Receivables: Credit Risk

Reinsurance receivables, in particular those due the ceded catastrophe risk, constitute the greatest portion of the credit risk faced by an insurance company. Seabury ERM breaks down reinsurance receivables by the reinsurer and then proxies the credit risk of the receivables by the credit risk of a portfolio of risky bonds. The credit rating of the reinsurer gets assigned to the bond, and the amount of each receivable becomes the bond face value.

3.2.2 Integration and Risk Aggregation

ERM is a fully integrated model. ERM does not make any assumptions about a company's exposure to different risk categories or about correlations between them; on the contrary, these characteristics arise logically from ERM analysis. ERM employs a bottom-up approach to risk aggregation—the only proper way to handle correlations between risks of a dynamic company in a changing environment. The top-down approach is based upon the presumption that the total corporate structure could be broken down to a number of high level business segments and the total risk could be derived from the correlation structure between the segments. While this method is easy to implement, it has some major deficiencies:

1. There is usually no statistically sound way to assess these high level correlations. As a result, companies have to rely on expert opinion to come up with the estimate for the correlations between segments. Expert opinions may be irrelevant in a dynamic market environment.
2. Business segments are not static; their composition may transform over time. While weights of the separate components within the segment

change, correlations between the segments change as well, and this, in turn, may lead to the risk miscalculation.

In contrast, ERM parameters come from current market data and the correlations are measured starting from the lowest level (such as individual positions or lines of business (LOBs)) and then expanding up to any desired level: for example, to wholly owned subsidiaries or sister companies.

3.2.3 Correlation Structure and Calibration

All ERM risk factors exhibit correlations. It is easy to observe that the yields of different maturities move in tandem, or that the equity returns of different sectors are usually highly correlated. Yet it is often the case that macroeconomic factors drive the premium rates as well as frequency and severity of the insured events, just like they affect the default probabilities in the credit models. For example, an increase in the interest rates will result in larger profits from invested premiums, which, in competitive markets, will apply downward pressure to the premium rates. Recessions lead to increased frequency of workers' compensation claims and increased average age of the insured cars.

ERM risk factors and their correlation structure are estimated from the latest market data and both public and proprietary insurance data sources. The long-term horizon required by the nature of insurance business and the distinct differences between investment instruments and insurance liabilities present a number of unique problems for the proper calibration of the correlation structure. Among them:

- **Stability.** The random noise inherent to large correlation matrices will, in effect, get amplified over a long forecasting period and, if not filtered out, render the risk statistics unreliable. This is especially true if any optimization of a business structure is attempted while making use of a raw correlation matrix.
- **Correlation between assets and liabilities.** Large correlation matrices commonly employed on the investment side require long time series for estimation. The large required number of data points does not present a problem for financial time series which are generally available on a daily basis. On the other hand, the natural time unit for insurance data is one year, and it would be practically impossible to collect a sufficient history of premium and loss data without incurring complicated seasonal effects.

As a result, there is no feasible way to represent both investment assets and insurance liabilities in the form of a “grand” correlation matrix.

In order to assure stability and the statistical significance of the estimates, Seabury ERM applies various calibration techniques, such as principal component analysis (PCA) and regression analysis. ERM employs the latest techniques of random matrix theory developed in physical science (Sengupta and Mitra, 1999) in order to discriminate between significant components and noise. PCA drastically reduces the “effective” dimensionality, which allows establishing the dependence of the insurance risk factors on the macroeconomic environment through regression analysis (0).

3.2.4 “Stable” Simulation

A single set of risk measures may satisfy a regulatory requirement but it is unlikely to give a complete picture of a company’s risk. Risk managers will want to know the sensitivity of the model’s output relative to changes in input assumptions such as expected equity returns and premium rates. They will also want to investigate different scenarios such as a removal of a specific holding or instrument class, or an addition of a line of business.

Simulation models always present a problem in this regard due to the inherent variability of simulation results between any two runs. When input assumptions change, the simulation variability may easily mask the sought changes in the output. As a result, the sensitivity and directional analysis can be rendered questionable unless an extremely large number of simulations are performed in each run. Even the discrete changes of scenario analysis become problematic. Thus, the requirement that “small changes in the simulation input always lead to small changes in the output” becomes of critical importance. We refer to such quality of a simulation model as “stable simulation.”

Seabury ERM solves this problem by employing a patented “seeding” technique. Each and every variable receives its own random or quasi-random number generator, which is uniquely seeded (initiated) depending on the financial instrument or risk factor the variable describes (0). The end result is a “stable simulation” framework that always allows two different runs to be directly compared, as long as the number of simulations does not change.

3.3 Risk Reporting

An enterprise risk model should provide a comprehensive set of reports that enable risk managers to see different aspects of a company's risk. Risk reporting usually occurs at different levels within a company including the corporate level, subsidiary level and business unit level. It is crucial that business units and subsidiaries use the same methods and follow the same guidelines to produce a standardized set of risk reports across the corporation.

The risk reporting capabilities possessed by ERM allow its users to analyze many different aspects of risk. ERM provides risk reports by risk category or business line; it can combine portfolios, individual companies, or company groups. This integrated approach to risk measurement enables rapid creation of new reports and customization of the existing reports upon request.

Banks and securities firms calculate risk at a short-term (3 to 10 trading days) horizon appropriate for trading operations. They traditionally report risk in terms of VaR, but, often, a simplifying assumption of a normal distribution is made, and VaR is calculated in terms of standard deviation. For insurance companies, on the other hand, one year is a natural reporting horizon. Insurance operations, as a rule, are prone to rare but significant losses which manifest themselves in the highly non-normal, fat-tailed P&L distributions. Unlike the standard deviation, the properly calculated VaR-based metrics focus on the tail of the P&L distribution—on rare events that may threaten the very solvency of the company. These considerations are endorsed by regulatory authorities. The proposed Solvency II regulations require insurance companies to calculate VaR and similar risk measures on a longer-term basis of a one-year horizon. Solvency II specifically encourages the use of Tail VaR (TVaR, Expected Shortfall) (see Committee of European Insurance and Occupational Pensions Supervisors, 2005, p. 105) defined as the expected amount of loss when the loss exceeds VaR. TVaR is conceptually close to measuring the risk in terms of the value of a hypothetical put option that would be required to completely hedge the losses over a certain threshold. This metric possesses a number of desired qualities (Artzner et al., 1999) and has become increasingly popular in risk management.

The standard ERM risk reports include profit-and-loss (P&L) distribution and such risk measures as:

- **Standard deviation**—by business unit and risk category.
- **VaR**—by business unit and risk category.

- **Expected Shortfall**, also known as **Tail VaR (TVaR)**—by business unit and risk category.
- **Marginal VaR and TVaR**—by business unit. Marginal VaR refers to the change in the company's VaR when a business unit is added as a whole.
- **Incremental VaR (IVaR) and TVaR**—by business unit. IVaR measures the change relative to a gradual expansion of the unit.
- **Downside probabilities**—at the corporate level. Downside probability is defined as probability of losing a certain fraction of the net worth.

VaR-based metrics measure the risk of rare events, such as catastrophes. If there happen to be just a few of scenarios with such events, these tail risk measures may not be estimated reliably. Seabury ERM improves the convergence of the VaR metrics with help of importance sampling and statistically robust L-estimation procedures (Borovkov et al., 1996 and Mausser, 2001) (see Appendix D for details of these techniques).

In addition to various risk measures, ERM calculates economic capital (EC) required to support the solvency at a given probability level. The EC is allocated to each business unit in proportion to their IVaR, as supported by modern financial science. Finally, to facilitate performance measurement, RAROC is calculated for each business segment.

3.4 RAROC and Capital Allocation

After the company's overall risk has been assessed along with the risk contributions from each of the principal risk categories, management will want to assess the performance of each of its business segments. The relative profitability of different business segments could be gauged from the RAROC estimates calculated for each segment.

The authors believe that individual business return performance should be based on stand-alone RAROC rather than providing an allocated diversification benefit to each business. We believe that any benefit that accrues to a business activity resulting from the company's portfolio structure should be credited to a general corporate account rather than to the individual business activity that had nothing to do with creating the benefit.

3.4.1 LOB RAROC

ERM breaks the risk, capital and RAROC for each business activity or LOB down into reserve risk and underwriting risk components. The benefit of breaking the LOB performance down into reserve (past business) and underwriting (new business) parts is that it allows management to evaluate the effectiveness of its pricing policy on past as well as future business. The total RAROC will not provide this information.

RAROC on the reserve risk represents the return for past business that has already been accepted and will include all the reserve strengthening that has occurred on that business to date:

$$\begin{aligned} \text{Reserve RAROC} &= \\ &= \frac{\text{PV}(\text{Duration matched interest on insurance reserves}) + \text{Interest on EC}}{\text{PV}(\text{EC})}, \end{aligned} \quad (3.4.1)$$

where economic capital EC refers to the stand-alone risk of an LOB.

Underwriting risk RAROC represents the expected return that the company should obtain for the premium that it will collect in the next period—the next 12 months. Management will want to input into the ERM model all of the risk factors that it believes may be experienced in the next period, i.e., premium rate increases or declines, premium volume, expense, economic and investment rate factors, etc. In this way, management will learn what type of return it should expect to make in the next period given its current pricing strategy.

$$\begin{aligned} \text{Underwriting RAROC} &= \\ &= \frac{\text{PV}(\text{Premiums}) - \text{PV}(\text{Expected Losses}) - \text{PV}(\text{Expenses}) + \text{Interest on EC}}{\text{PV}(\text{EC})}, \end{aligned} \quad (3.4.2)$$

where $Premiums$ —to be collected in the next period; $Expected Losses$ (losses and ALAE)—for all future periods while reserve is still active; $Expenses$ (ULAE and overhead)—for all future periods while reserve is still active.

3.4.2 Investment RAROC

Investment RAROC measures the effectiveness of the investment policy put into practice by the investment department:

$$\text{Investment RAROC} = \frac{(R_p - P_t) \cdot MV_p}{\text{PV}(\text{EC})}, \quad (3.4.3)$$

where economic capital EC refers to the stand alone risk of the investment segment, MV_p is the market value of the investment portfolio, R_p – return on the portfolio, P_t – internal transfer rate, i.e., the rate the investment department uses to borrow money from the insurance group. Investment department then credits the insurance group with the risk-free duration matched (to duration of reserves) interest rate on funds borrowed. The internal transfer rate may also be adjusted to reflect the firm’s external borrowing.

3.4.3 Total Company RAROC

The total RAROC sums up the reserve risk, underwriting risk and the investment risk. Overhead is allocated to each business unit in accordance with actual usage.

$$\begin{aligned} \text{Combined RAROC} &= \\ &= \frac{\text{Income}(\text{Underwriting}) + \text{Income}(\text{Reserve}) + \text{Income}(\text{Investments})}{EC(\text{Company})}. \end{aligned} \quad (3.4.4)$$

Total company RAROC becomes particularly important when the company’s performance is evaluated relative to a peer group, or compared to the average industry performance.

3.4.4 Other Performance Measures

It is important to note that RAROC performance should be a significant criterion of financial performance measurement, but it is not the only important criteria. RAROC, by itself, will not inform management of the shareholder value contribution of a business activity. An important driver of shareholder value is earnings growth, which is not measured in RAROC. Also, shareholder value measurement measures the return on a company’s market value, not on its economic capital; and the required return (RR, the discount factor applied to the company’s earnings) that is embedded in the shareholder value formula includes an adjustment for systematic risk, while RAROC does not. We can express the shareholder value in at least two different ways:

$$\frac{\text{Market}}{\text{Book}} = \frac{ROE}{RR - Growth}, \quad (3.4.5)$$

or:

$$\text{Value} = \frac{\text{Earnings}}{RR - Growth}. \quad (3.4.6)$$

The reader will observe that there is a significant difference between these two methods of measuring financial performance. A one-period model for RAROC puts all the emphasis on earnings that are collected for, at most, one more year. Investment income in RAROC is present valued at the risk-free rate of interest (or some other transfer rate). Future growth of earnings possibilities is not considered under RAROC. For this reason, the authors caution that, while RAROC is an indispensable piece of information for assessing financial performance, it is not the only earning performance measure that should be considered. However, the purpose of this paper is to present and discuss a unified risk theory for measuring insurance company risk and capital requirements. The authors acknowledge that there is not agreement between RAROC, as defined here, and shareholder value measurement. However, we do believe that RAROC, as defined, does reveal the relative financial return on risk-adjusted capital among a company's business activities or lines of business. We reserve the discussion of shareholder value unification with that of RAROC for a future paper.

3.4.5 Capital Allocation

Capital allocation is another major subject where ERM may assist the management in making strategic decisions. The authors have used the IVaR of a business line as opposed to the stand-alone business risk to allocate capital on what we call an "economic" basis. By economic basis of capital allocation (as contrasted with that of stand-alone business risk allocation), we mean that management should know where its capital is actually being allocated to support risk in contrast to how we advocate that individual business performance should be measured on a stand-alone basis. This is in keeping with our prior statement that individual business performance should not receive a credit that is not the result of its own activity. The use of IVaR for economic capital allocation purposes ensures that capital will be allocated on an economic basis to each business activity in proportion to the risk that each activity contributes to the total company risk. The stand-alone risk allocation of capital to business activities would not be appropriate for understanding the economic allocation since the sum of the individual stand-alone risks will be larger than the total risk of the enterprise owing to the correlations between the business sectors. As a result, the total required enterprise capital may be miscalculated.

4. ERM as a Strategy-making Tool: Case Study

A versatile risk management platform should also provide management with a tool to assess the effectiveness of its business strategies. ERM provides this capability by having a robust scenario-generating capacity that allows its users to test the risk and profitability of each business strategy. We identify the following classes of scenario:

- **Changes in market environment**—analyzes the impact of changes in market conditions, like interest rates, return on stock market, exchange rates...
- **Changes in business environment**—analyzes the impact of changes in premium rates and their volatilities, expansion/contraction of the business, adding/removing lines of business...
- **Changes in the investment portfolio**—analyzes the impact of changes in the investment portfolio structure, including modifications of portfolio composition, adding/removing specific holdings or instrument classes...
- **Mergers and acquisitions**—potential effect that acquisitions or mergers may have on the individual company or merged group of companies.

4.1 Case Study—Cargo Inc.

Cargo Inc. (Cargo) is a multi-line property and casualty company that writes business on a nationwide basis. Cargo is a real company but the authors have disguised its name and its numbers. Only publicly available information has been used to assess the operations of Cargo.

1. The analysis was compiled from the following sources of information:
 - a. 2004 Regulatory Statement including Schedules D and P from Highline Data
 - b. 2004 Annual Statements and 10Ks and 10Qs
 - c. Bloomberg and Dow Jones for financial data
 - d. CMS BondEdge for key rate durations and cash flow projections for all ABS & MBS securities

- e. Only publicly available sources (retail or otherwise) of information were used—no company contact
2. The analysis consists of a base case and five scenarios
 3. The analysis will include insurance risk, interest rate risk, equity risk, credit risk and foreign exchange risk
 4. The analysis will not include
 - a. Catastrophe risk; unless captured in the last 10 years from Schedule P
 - b. Operating risk
 - c. Prior year losses (i.e., losses unpaid older than 10 years)

The following two tables summarize the investment portfolio of Cargo and its business lines:

TABLE 4
Cargo's Investment Portfolio

Asset Class	Market Value	Composition
Sovereign Bonds	670,000,000	12%
Municipal Bonds	83,000,000	2%
Corporate Bonds	2,166,000,000	40%
ABS/MBS	2,210,000,000	40%
Common Stock	329,000,000	6%
Preferred Stock	0	0%
Total	5,458,000,000	100%

TABLE 5
Cargo's Lines of Business

LOB	P&C Business Line	Premium
A	A. Homeowners/Farmowners	266,985
B	B. Private Passenger Auto Liability/Medical	92,381
C	C. Commercial Auto/Truck Liability/Medical	102,181
D	D. Workers' Compensation	93,057
E	E. Commercial Multiple Peril	400,575
F	F1. Medical Malpractice—Occurrence	53,198
FC	F2. Medical Malpractice—Claims-Made	17,836
G	G. Special Liability (Ocean Marine, Aircraft (All Perils), Boiler And Machinery	78,070
H	H1. Other Liability—Occurrence	339,150
HC	H2. Other Liability—Claims-Made	65,150
M	M. International	0
N	N. Reinsurance A—Non-proportional Assumed Property	0
O	O. Reinsurance B—Non-proportional Assumed Liability	0
P	P. Reinsurance C—Non-proportional Assumed Financial Lines	0
R	R1. Products Liability—Occurrence	25,582
RC	R2. Products Liability—Claims-Made	145
I	I. Special Property (Fire, Allied Lines, Inland Marine, Earthquake, Glass, Burglary And Theft)	490,777
J	J. Auto Physical Damage	88,196
K	K. Fidelity/Surety	11,823
L	L. Other (Including Credit, Accident and Health)	39,461
SS	S. Financial Guaranty/Mortgage Guaranty	0
Total		2,164,563

Our primary goals with this case study are to:

1. Determine the company's VaR at the 99 percent confidence level
2. Determine how much risk is contributed by the five categories of risk
3. Determine how this capital needs to be allocated to Cargo's different business activities
4. Determine the company's downside risk
5. Determine the impact to the company of various economic and stress scenarios

Seabury set up a base case assessment of Cargo's risk with alternative scenarios to assess the impact that these particular scenarios would have on the company's VaR:

Base Case: Steady-State

- Based on Cargo's 2004 Regulatory Statement
- Assumes Cargo will write the same amount of new premium in 2005 as it wrote in 2004, i.e., same operating ratio of 1.3
- Assumes that premium rates will remain the same in 2005 as in 2004 for all lines
- Assumes the premium volatility is zero
- Assumes a 7 percent return on all U.S. equity sectors.

Scenario 1:

- Same as Base Case, except that Scenario 1 assumes that premium rates on the biggest line, commercial multiple peril, decline 5 percent with 10 percent premium rate volatility.

Scenario 2:

- Same as Base Case, except that Scenario 2 assumes that U.S. yield curve shifts upward 100 bps.
- *Comments:* We believe that any company should be aware of its sensitivity to yield curve shifts. A 100 basis point shift is not terribly significant and management should anticipate moves of at least this magnitude.

Scenario 3:

- Same as Base Case, except that Scenario 3 assumes that the return on U.S. stock market is a negative 10 percent.
- *Comments:* Companies will want to understand the sensitivity of their equity portfolios to equity market movements. We believe that there is lots of precedent for equity market returns of minus 10 percent and that this is a realistic scenario for management to observe.

Scenario 4:

- Combination of Scenarios 1, 2 and 3. Premium rates for commercial multiple peril decline 5 percent with 10 percent premium rate volatility, U.S. yield curve shifts upward 100 basis points, the return on U.S. stock market is a negative 10 percent.
- *Comments:* Risks are typically correlated in one way or the other. The interactions of these risks can be quite complicated and the complete distribution of their mutual integration needs to be understood in order to

assess the potential impact of complex economic events. The authors believe that Scenario 4 is a probable event for which Cargo should be prepared.

Scenario 5:

- Premium rates for all LOBs decline 15 percent with 20 percent premium rate volatility, U.S. yield curve shifts upward 300 bps, the return on U.S. stock market is a negative 10 percent.
- *Comments:* This scenario represents an extreme stress-test to which the authors subjected Cargo. We believe that there is precedent in the historical markets for the potential realization of this event and therefore something that Cargo's management should consider. For example, between September of 1980 and September of 1981, the yield on 10-year Treasury obligations rose by 398 basis points while equity markets dropped 7 percent. There have been openly speculative, yet sober, discussions occurring in prominent economic circles about a potential "hard-landing" for the U.S. economy. This scenario has the U.S. twin deficits tempering the amount of U.S. Treasuries foreigners are willing to hold; interest rates spike, equity markets tumble, and a recession ensues in which the prices of goods and services (along with insurance prices) also fall. The authors submit this scenario as an example of an extreme stress test and leave the weighting of its validity to Cargo's management.

The base case is intended to evaluate the risk and return of Cargo under conditions as close to the business climate as has prevailed for the last 10 years as possible. In other words, the base case should reflect management's view of the firm's required capital in the ordinary course, or steady-state. Each of the subsequent scenarios is intended to simulate a condition that management believes has a probability of occurring. How high of a probability is a matter that will be discussed later.

The objectives of the ERM analysis should be to determine if the company has: 1) enough economic capital to support its business activities at a specified level of statistical confidence through a variety of adverse economic scenarios; 2) to develop an understanding of the company's downside risk, i.e., the chance probability of losing a certain amount of its capital in any one year; 3) to ensure that capital is allocated to business activities in proportion to their respective risk contributions; and 4) to ensure that the rate of return on capital for each business activity is understood.

Management should know the answer to each one of these questions for each business strategy that it undertakes. Each business strategy should also be stress-tested. In this way, the risk and return of different strategies can be compared.

4.2 Four Principal Drivers of Economic Capital

The authors believe that there are four principal drivers of the amount of economic capital that companies should carry:

1. Capital in the ordinary course, i.e., steady-state capital

The amount of capital that is required given a company's VaR at a specified level of statistical confidence. The statistical confidence results from the model's data parameterization period. Steady-state capital does not capture rare market events. To capture this issue, management needs to perform stress-tests.

2. Operational risk capital

The authors define operational risk capital as the capital that is required to support losses that result from a company's failed internal processes, such as losses that result from a failure to draft contracts tightly; losses that result from systems failures, etc.

3. Capital for adverse market movements—stress-tests

The amount of capital that is required to sustain a company through severe market dislocations and volatility, also known as market events. We may know these events by names such as: October 1987, The Asian Crisis of 1997, or the equity market crash of 2001. The point is that these events are often not represented in the historical data parameters of the enterprise risk model—any model has this limitation. As such, these events need to be imposed on the model in order to learn what their impact would be on the company given the particular structure of its assets and liabilities.

4. Catastrophe risk capital

The amount of capital that is required by a company to safely absorb losses that result from its catastrophe exposures.

4.3 Risks Omitted from this Case Study

It is important to say at the very beginning that this case study will not attempt to evaluate the risk capital that will be required to support Cargo's operational or catastrophe risk (other than what has been captured in the last 10 years of the paid loss triangles). The principal purpose of this exercise is to reveal the dynamics of measuring integrated risk with those fields of information that are available to us on a robust basis from publicly available sources. This includes: interest rate risk, credit risk, foreign exchange, and insurance risk (non-catastrophe). Reliable figures on operational and catastrophe risk are not readily available from any source other than the company itself. Therefore, the economic capital assessments and allocations that are made in this exercise may appear low due to the omission of these two factors. Seabury's ERM engine does possess the ability to readily integrate and aggregate a company's operating risk and catastrophe distributions with that of the rest of its risk so that a total solution is realized. *The principal issue that we are attempting to illustrate in this case study is how economic capital is impacted within an integrated framework, from a base case realization to that of various stress-tests.*

4.4 Operational Risk Capital

The difficulty of quantifying operational risk is that there are no public databases that record and classify these events such as there are for most of the other categories of risk—interest rate risk, credit risk, paid losses, etc. As such, the first thing that any company will discover as it implements an integrated enterprise risk management system is that the economic capital requirement will often be significantly less than the company's actual capital level (nominal capital). The difference between companies' nominal capital and economic capital should generally be attributed to three sources of risk: 1) operational risk; 2) catastrophe risk; and 3) market event risk.

We have identified catastrophe risk as well as what the grab-bag of operational risks may consist. Market events, and the capital required to support them, can be described as those events that occur in the market whose expectation may fall outside the model's historical data parameters. For example, Seabury's ERM is parameterized with 10 years of data for both assets and liabilities for which we obtain annual updates. This means that Seabury's ERM utilizes historical probabilities of the events that have occurred over the last 10 years. The resulting VaR represents a company's risk given that events during this historical time period are as equally likely to occur over the next year as they were over any one of the previous 10 years.

Should we assume that Cargo has sufficient capital if its nominal capital exceeds its economic capital? To answer this, we would have to ask ourselves if the financial and insurance markets might experience an event during the next year that was without precedent in the last 10 years. If so, we would want to include a scenario that reflected the potential for that event; be it a market movement like that of October of 1987, the Asian Crisis of 1997 or the equity market crash of 2001. This is why we create scenario analysis and stress-tests. This is also why companies' nominal capital will generally exceed, by some significant margin, their economic capital requirements as measured within an ERM context. The drafters of risk-based capital for such regulatory initiatives as Basel II and Solvency II have debated and openly speculated as to which multiplier of the 99 percent VaR economic capital should be imposed on the insurance and the banking industries to account for this limitation. Since this risk parameter is unknowable, risk managers often prefer to implement scenario analysis and stress testing to arrive at their best estimate of required capital. For this exercise on Cargo, we are not including calculating the capital requirements for operational or catastrophe risk requirements, only that risk resulting from steady-state conditions and that risk resulting from various scenarios and market events.

4.5 Cargo's Risk

In TABLE 6 below, we display the results of Cargo's risks by category and scenario. Risks are displayed in dollars (millions).

TABLE 6

	Base Case	Scenario 1	Scenario 2	Scenario 3	Scenario 4	Scenario 5
Reserve Risk	505	505	505	505	505	505
Underwriting Risk	132	161	132	132	161	434
ALM (Interest) Risk	236	236	220	236	220	189
Equity Risk	108	108	108	93	93	93
Forex Risk	5	5	5	5	5	5
Credit Risk	107	107	107	133	133	133
Total	1,093	1,121	1,076	1,103	1,115	1,359
Diversification Benefits	495	520	465	506	503	545
Required Economic Capital	598	602	611	598	613	813

To be able to make any assessments about the company's capital adequacy, we need to examine the company's expected year-end surplus in comparison to its required level of capital at the 99 percent confidence level. We can observe this in TABLE below.

TABLE 7

	Base Case	Scenario 1	Scenario 2	Scenario 3	Scenario 4	Scenario 5
Expected Year-End Surplus	2,284	2,264	2,166	2,238	2,098	1,493
Expected Year-End Income	637	616	462.5	521	394	-312
Economic Capital at 99% Level	598	602	611	598	613	813
Prob. of losing 10% or more from beginning of year surplus	0.1%	0.2%	1.4%	0.3%	2.2%	68.5%
Prob. of losing 10% or more from expected end of year surplus	17.4%	18.1%	19.8%	17.9%	20.4%	30.8%
Prob. of default	0%	0%	0%	0%	0%	0%
Operating ratio from beginning of year surplus	1.26	1.26	1.23	1.26	1.23	1.18

TABLE reveals the impact of each scenario to several critical measures of economic capital and income. For the most part, it appears that Cargo has enough economic capital to sustain it through all of the scenarios, although Scenario 5 has a serious impact.

TABLE includes some entries that have not yet been defined:

Expected year-end surplus: This represents the amount of surplus that ERM predicts that Cargo will have at the end of the year and includes adjustments that result from: retained earnings (pre tax), capital gains/losses, unrealized gains/losses and changes to reserves.

Expected year-end income: This represents the earnings that ERM predicts Cargo will be able to recognize at year end and includes adjustments resulting from: capital gains/losses, unrealized gains/losses and changes to reserves. Note: The definition of income in the ERM framework is equivalent to that of net worth (i.e., impact to surplus), not distributable earnings.

Probability of losing 10 percent or more from beginning-of-year/end-of-year surplus: This represents the probability that the company will lose at least 10 percent or more of its capital from its beginning-of-year/end-of-year surplus level. This is also known as the company's downside risk. ERM assesses this probability for any percent of the firm's surplus.

Key insights from ERM analysis of Cargo:

- **Base case:** As presented, Cargo has nearly four times the economic capital that it requires.
- **Scenario 1:** There is no increase in reserve risk as new premium rate declines will not impact reserve risk. Underwriting risk does increase as was expected given that premium rates in Cargo's biggest line have decreased by 5 percent with an increase in the premium rate volatility to 10 percent. All other risks are unaffected.
- **Scenario 2:** The value of the fixed income portfolio declined due to the rise in interest rates. At the same time, the present value of its liabilities also declined reducing total ALM risk. Overall risk would have declined in Scenario 2 from that of the base case except that this change in Cargo's portfolio structure reduced its diversification benefit leaving the required economic capital almost unchanged.
- **Scenario 3:** As expected, the equity risk declined owing to the lower base rate return. The credit risk increased owing to the integral role that equity value contributes to solvency in a contingent claims valuation framework.
- **Scenario 4:** Company loses approximately 38 percent of its income from that of the base case due to all three effects, but due to diversification, this combined risk scenario contributes little risk.

- **Scenario 5:** Due to the magnitude of the loss in income from the severity of these events, we now observe negative income and a significant increase in company risk. Probability of losing 20 percent of the initial surplus is now 45.4 percent.

What can the CEO of Cargo say to his board of directors about their company's overall risk and financial performance based on this analysis?

- That the company appears to be well capitalized in the base case and first four economic scenarios.
- That in Scenario 5, the company would have to increase economic capital by about 36 percent to maintain the same level of downside risk as exists in the base case today.

Impact to income:

- In Scenario 2, income is reduced 27 percent.
- In Scenario 3, income is reduced 18 percent from the base case due to a lowering of equity market returns and an increase in the discount rate that lowers the present value of the fixed income portfolio.
- In Scenario 4, income is off 38 percent due to the combined effect of Scenarios 1, 2, and 3.
- In Scenario 5, income is a negative \$312 million.
- Pre-tax RAROC for company appears to be adequate in all but Scenario 5.
- The company's earnings appear to be quite sensitive to yield curve shifts. Cargo may want to explore hedging opportunities and/or better ALM matching.

What the CEO actually says to his board will depend, in large part, on how large a probability his advisors will assign to Scenarios 4 and 5, since this is where both economic capital and earnings get hit the hardest. The CEO will have to review how significant an event it would be for Cargo to lose 20 percent of its capital since the probability of this occurrence increases to 45.4 percent in Scenario 5. To be clear, the analysis is not saying that there is a 45.4 percent chance that Cargo will lose 20 percent

or more of its capital, it is saying that this is only likely in the event of Scenario 5—and Scenario 5 may only have a 5 percent chance of occurring. There is no way of attributing the likelihood of a market “event” unless that event is contained in the model’s data history. Assigning weights to low probability scenarios is beyond the scope of what an ERM system can provide if the events are not in the model’s historical parameters. However, an integrated ERM can provide the user with a good assessment of the likely impact if the event does occur.

4.6 Why are Scenarios and Stress Testing Important?

Stress testing is a process of identifying, often through Monte Carlo Simulations, the response of an asset or liability portfolio (or both), to a variety of types of financial distress. The purpose of the stress test is to evaluate how the organization whose portfolio is being tested would fare under a specific set of adverse market conditions. In the absence of being able to identify how many times 99 percent VaR capital should be imposed, risk managers will often identify specific market crises as being the events for which they want their institution to be prepared. These events, and their demonstrated ability to survive them without impairment, become a part of what they represent to the world as their statement of solidity. For this reason, ERM has been provided with a robust scenario-generating capability. As a result, the ERM user can use an ERM system in two ways: 1) to observe their economic capital requirements at a desired level of VaR confidence assuming the business and economic climate does not deviate substantially from steady-state; or 2) the user can stress test market events that may have low expectations of occurrence—perhaps as low as once in 50 years for financial risk or once in 500 years for catastrophe risk. Many companies have found it useful to think in these terms rather than establishing a multiple of steady-state capital as the desirable level of capital to hold.

Cargo’s management team will want to run a large variety of single and stacked (simultaneous) scenarios that reflect their judgment about potential market events. They will then want to position their company to be able to withstand a level of risk that is in concert with their and the board’s collective judgment about acceptable risk. To be able to do this, Cargo will need an integrated ERM that is fast enough to be able to provide results in real time and simple enough that complex risk interactions are revealed in easy-to-digest output. Cargo’s management may want to increase capitalization by about 36 percent if it wants to maintain the same downside risk probabilities in Scenario 5, as it currently enjoys in the base case. The countervailing judgment would have to be: how disadvantageous would this level of capitalization be to the company’s current market competitiveness? This is the balance that management must strike.

4.7 Capital Allocation and RAROC at Cargo

ERM allocates economic and stand-alone capital allocations to Cargo's regulatory operating segments. Several lines require negative capital due to the diversification effect. Nominal capital has been allocated in proportion to that of economic capital and overhead has been allocated to each LOB in proportion to its premium volume (since specific overhead allocations are not available to the authors). RAROC for individual lines has been calculated based on stand-alone risk. Most of the individual segments appear to be making adequate RAROC values.

TABLE 8

Segment		Standalone Risk	Economic Capital (Allocated Risk)	Company Capital Allocated	Income	RAROC
Investment		469,976	219,334	838,735	169,829	36%
A. Homeowners/Farmowners	Old	26,012	-7,153	-27,354	2,573	
	New	10,551	-1,230	-4,702	43,393	
	Total	31,938	-8,383	-32,056	45,966	144%
B. Private Passenger Auto Liability/Medical	Old	37,987	20,942	80,083	2,409	
	New	10,816	5,269	20,148	-7,271	
	Total	47,592	26,211	100,231	-4,862	-10%
C. Commercial Auto/Truck Liability/Medical	Old	43,694	7,465	28,545	4,336	
	New	18,712	1,355	5,180	-1,459	
	Total	58,940	8,819	33,725	2,877	5%
D. Workers' Compensation	Old	35,817	-17,208	-65,804	15,342	
	New	3,691	-1,225	-4,684	18,173	
	Total	37,620	-18,433	-70,488	33,515	89%
E. Commercial Multiple Peril	Old	80,984	35,007	133,867	12,778	
	New	26,321	1,809	6,918	66,052	
	Total	100,185	36,816	140,785	78,830	79%
F1. Medical Malpractice - Occurrence	Old	11,866	3,514	13,438	2,396	
	New	3,292	-199	-761	25,550	
	Total	13,413	3,316	12,678	27,945	208%
F2. Medical Malpractice - Claims-Made	Old	211,080	132,431	506,419	6,166	
	New	39,511	11,566	44,228	-22,862	
	Total	228,840	143,997	550,646	-16,696	-7%
G. Special Liability	Old	16,186	-4,481	-17,134	1,999	
	New	9,012	-305	-1,167	15,072	
	Total	18,415	-4,786	-18,301	17,071	93%
H1. Other Liability - Occurrence	Old	146,388	85,918	328,550	20,586	
	New	45,623	19,676	75,240	52,331	
	Total	182,756	105,593	403,789	72,917	40%

Segment		Standalone Risk	Economic Capital (Allocated Risk)	Company Capital Allocated	Income	RAROC
H2. Other Liability - Claims-Made	Old	264,846	64,052	244,933	9,961	
	New	37,151	2,709	10,357	-5,413	
	Total	285,481	66,760	255,291	4,548	2%
R1. Products Liability – Occurrence	Old	40,745	4,271	16,334	2,728	
	New	258.5	60	229	17,727	
	Total	40,752	4,331	16,562	20,454	50%
2yr. Combined Business	Old	87,989	12,746	48,742	1,911	
	New	109,456	1,002	3,831	182,240	
	Total	127,086	13,748	52,572	184,151	145%
Total	Total	597,322	597,322	2,284,168	636,542	106.6%

4.8 Segment Level RAROC

At this level of focus we are better able to make certain distinctions. Cargo considers that its premiere segment is “specialty.” This may be true in terms of market share at acceptable rates, yet it makes a superior RAROC on both commercial and personal lines. Yet, we also see that commercial and personal lines together only comprise about 7.5 percent of the entire company’s economic capital compared to 56 percent for that of specialty.

TABLE 9

Segment		Standalone Risk	Economic Capital (Allocated Risk)	Company Capital Allocated	Income	RAROC
Investment		469,976	219,334	838,735	169,829	36%
Personal Lines	Old	31,839	13,789	52,729	4,982	
	New	14,238	4,039	15,446	36,123	
	Total	42,344	17,828	68,175	41,104	97%
Commercial Lines	Old	121,023	25,264	96,608	32,456	
	New	39,928	1,939	7,414	82,766	
	Total	154,211	27,202	104,022	115,221	75%
Specialty	Old	470,302	298,451	1,141,280	45,744	
	New	130,420	34,507	131,956	264,644	
	Total	506,674	332,958	1,273,237	310,388	61%
Total	Total	597,322	597,322	2,284,168	636,542	106.6%

4.9 Final Conclusions on the Analysis

The predicate of this analysis was to determine:

1. Cargo's VaR at the 99 percent confidence level
2. How much risk is contributed by the five categories of risk
3. How this capital needs to be allocated to Cargo's different business activities
4. The company's downside risk
5. The impact to the company of various economic and stress scenarios.

We further pointed out that we did not have access to certain key types of information that would allow us to capture Cargo's operational and catastrophe risk. As a result, we limited the scope of our investigation to Cargo's economic capital requirements exclusive of these two risks. We further pointed out that the most critical part of our investigation was to examine:

How economic capital is impacted within an integrated framework, from a base case expectation to that of various stress-test levels.

We believe that we accomplished this objective. We learned that Cargo's surplus appears to be well structured against all but the most severe of the scenarios that we tested and that it appears to have more than sufficient capital in steady-state. From this foundation, it would be a relatively simple task to aggregate Cargo's catastrophe risk with the rest of its risk categories (except that of operating risk) to arrive at the total economic capital required. The required operational risk capital will be the difference between the firm's nominal capital and the combination of its economic and catastrophe risk capital. To assess if this level of operational risk capital is appropriate, we would want to benchmark Cargo with a peer group of companies taking the same integrated risk measurements for each of them as we have for Cargo. This task can be accomplished in the Seabury ERM.

5. Conclusion

Historically, leading financial institutions developed their own risk systems. Smaller companies generally had none. Today, an increasing number of companies outsource risk technology development to leverage the work that has already been put into developing standard risk analytics and to avoid costly mistakes. Choosing the right risk management system is an important investment decision with long-term implications. Risk managers should be clear about what they are looking for in a risk application, taking into consideration both current and projected business needs.

This document provides a brief overview of the methodology currently used by Seabury Analytic in our risk management applications. The methodology and algorithms used by ERM have been developed to be consistent with the best practices accepted by leading financial institutions. The models, assumptions, and techniques described in this document lay a solid methodological foundation for risk measurement in the insurance industry.

Risk management is not a precise science; there are a variety of methods and algorithms that could provide a good risk solution. Innovation, new business experience, or a change in regulations can encourage risk managers to revise and update risk methodologies. This implemented system possesses a great degree of flexibility: virtually every algorithm in ERM can be updated or replaced with a new one, as long as all modifications agree with the base rules utilized by the integrated ERM framework.

While every risk measurement system will require customization to meet specific user specifications, Seabury ERM includes all of the core capabilities that are required to meet best practice risk management and reporting requirements. But ERM is more than a compliance solution; it is also a business tool that allows users to quickly evaluate the effectiveness and profitability of each business strategy they may want to implement. Each business strategy can be assessed against a robust set of economic scenarios and stress tests, and the results are displayed in easy-to-digest metrics.

References

- Actuarial Standard of Practice No. 7, *Analysis of Life, Health, or Property/Casualty Insurer Cash Flows*, Actuarial Standards Board, Doc. No. 089 2002.
- Artzner, P., Delbaen, F., Eber, J., and Heath, D. 1999. Coherent measures of risk. *Mathematical Finance* 9 (3): 203–228.
- The Australian Prudential Regulation Authority. 2002, 2005. *Prudential Standard GPS 220 Risk Management*.
- Barnett, G., and Zehnwirth, B. 2000. Best estimates for reserves. *Proceedings of the Casualty Actuarial Society* LXXXVII: 245–321.
- Black, F., and Scholes, M. 1973. The pricing of options and corporate liabilities. *Journal of Political Economics* 81: 637–654.
- Borovkov, K., Dehling, H., and Pfeifer, D. 1996. On asymptotic behavior of weighted sample quantiles. *Mathematical Methods of Statistics* 5: 173–186.
- Bouchaud, J.-P., and Potters, M. 2003. *Theory of Financial Risk and Derivative Pricing* New York: Cambridge University Press.
- Committee of European Insurance and Occupational Pensions Supervisors. 2005. *Answers to the European Commission on the second wave of Calls for Advice in the framework of the Solvency II project*, CEIOPS-DOC-07/05.
- Doherty, N.A., and Garven, J.R. 1995. Insurance cycles: Interest rates and the capacity constraint model. *Journal of Business* 68(3): 383–404.
- England, P.D., and Verrall, R.J. 2002. Stochastic claims reserving in general insurance. *British Actuarial Journal* 8 (3): 443–518.
- Feldblum, S. 1995. Forecasting the future: Stochastic simulation and scenario testing, incorporating risk factors in dynamical financial analysis. 1995 CAS Discussion Paper Program: 151–177.
- The Financial Services Authority. 2005. *Integrated Prudential Sourcebook*. FSA Handbook (2005).

- Fung, H.-G., Lai, G.C., Patterson, G.A., and Witt, R.C. 1998. Underwriting cycles in property and liability insurance: An empirical analysis of industry and by-line data. *Journal of Risk and Insurance* 65(4): 539-561.
- Ghosh, S., and Puccia, M. 2004. *Property/Casualty Insurance Criteria: Capitalization*, Standard & Poor's.
- Grace, M.F., and Hotchkiss, J.L. 1995. External impacts on the property-liability insurance cycle. *Journal of Risk and Insurance* 62(4): 738-754.
- Gupton, G.M., Finger, C.C., and Bhatia, M. 1997. *CreditMetrics™—Technical Document*, New York: Morgan Guaranty Trust Co.
- Jarrow, R., and Turnbull, S. 1995. Pricing derivatives on financial securities subject to credit risk. *Journal of Finance* 50(1): 53-85.
- Jarrow, R. 1999. Estimating recovery rates and pseudo default probabilities implicit in debt and equity prices. *Financial Analysts Journal* 57(5): 75-92.
- Kalos, M.H., and Whitlock, P.A. 1986. *Monte Carlo Methods*. 1986. New York: John Wiley & Sons.
- Korobov, N.M. 1959. The approximate computation of multiple integrals. *Dokl. Akad. Nauk SSSR* 124: 1207-1210.
- Laloux, L., Cizeau, P., Bouchaud, J.-P., and Potters, M. 1999. Noise dressing of financial correlation matrices. *Physical Review Letters* 83 (7): 1467-1470.
- L'Ecuyer, P. 2004. Quasi-Monte Carlo methods in finance. *Proceedings of the 2004 Winter Simulation Conference*, IEEE Press: 1645-1655.
- Mausser, H. 2001. Calculating quantile-based risk analytics with L-estimators. *Algo Research Quarterly* 4(4): 33-47.
- Merton, R.C. 1973. Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4: 141-183.
- — — —. 1974. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance* 29: 449-470.

Mina, J., and Xiao, J.Y. 2001. *Return to RiskMetrics: The Evolution of a Standard*. New York: RiskMetrics Group.

RiskMetrics™—Technical Document, 4th ed. 1996. New York: Morgan Guaranty Trust Co.

Sengupta, A.M., and Mitra, P.P. 1999. Distributions of singular values for some random matrices. *Physical Review E* 80: 3389–3396.

Taylor, G. 2003. Chain ladder bias. *ASTIN Bulletin* 33(2): 313–330.

van Deventer, D.R., Imai, K., and Mesler, M. 2005. *Advanced Financial Risk Management*. Singapore: John Wiley & Sons, Asia.

Vasicek, O.A. 1977. An equilibrium characterization of the term structure. *Journal of Financial Economics* 5, 177–188.

Appendix A: Quasi Monte Carlo Simulation and Lattice Rules

A.1 Monte Carlo and Quasi Monte Carlo

The purpose of most stochastic simulations is to estimate the mathematical expectation of some cost function, in a wide sense. Since randomness in simulations is almost always generated from a sequence of i.i.d $U(0,1)$ (independent and identically distributed uniforms over the interval $[0,1)$) random variables, the mathematical expectation that we want to estimate can be expressed as the integral of a real-valued function f over the t -dimensional unit hypercube $[0,1)^t$,

$$\mu = \int_{[0,1)^t} f(\mathbf{u}) d\mathbf{u} \quad (\text{A.1.1})$$

For small t , numerical integration methods such as the Simpson rule, Gauss rule, etc., are available to approximate the integral (A.1.1). These methods quickly become impractical, however, as t increases beyond 4 or 5. For larger t , the usual estimator of μ is the average value of f over some point set $P_n = \{\mathbf{u}_0, \dots, \mathbf{u}_{n-1}\} \subset [0,1)^t$,

$$Q_n = \frac{1}{n} \sum_{i=0}^{n-1} f(\mathbf{u}_i). \quad (\text{A.1.2})$$

In the standard *Monte Carlo* (MC) simulation method, P_n is a set of n i.i.d. uniform random points over $[0,1)^t$. Then, Q_n is an unbiased estimator of μ with variance σ^2/n . When the variance is finite, the estimator Q_n converges to μ with the convergence rate $O(1/\sqrt{n})$.

The *Quasi-Monte Carlo* (QMC) method constructs the point set P_n more evenly distributed over $[0,1)^t$ than typical random points, in order to improve over the $O(1/\sqrt{n})$ convergence rate. The precise meaning of “more evenly” depends on how we measure uniformity, and this is usually done by defining a measure of *discrepancy* between the discrete distribution determined by the points of P_n and the uniform distribution over $[0,1)^t$. A *low-discrepancy point set* P_n is a point set for which the discrepancy measure is significantly smaller than that of a typical random point set.

The two main families of construction methods for low-discrepancy point sets in practice are the *digital nets* and the *integration lattices*. The former usually aim at constructing so-called (t, m, s) -nets. A *low-discrepancy sequence* is an infinite sequence of

points $P_\infty = \{\mathbf{u}_0, \mathbf{u}_1, \dots\}$ such that for all n (or for an infinite increasing sequence of values of n , e.g., each power of 2), the point set $P_n = \{\mathbf{u}_0, \dots, \mathbf{u}_{n-1}\}$ has low discrepancy. The number of points does not need to be set in advance, before running a simulation: one can always add additional points if the convergence is insufficient. Well known examples of (t, m, s) -nets are sequences constructed by Halton, Sobol', Faure, and Niederreiter. Unfortunately, the quality of these sequences deteriorates with the dimensionality t , while the complexity of their construction significantly increases; there are no reported examples of (t, m, s) -nets applied to problems with more than 400 dimensions.

In contrast to (t, m, s) -nets, the integration lattices are generally not extendable: they require the number of points n be set in advance (although, a construction of extendable integration lattices has recently been reported in the literature). If the desired convergence error has not been achieved with a given n , one needs to run the simulation all over again with a greater n . At the same time, the construction rules, as explained below, are very simple, and the integration lattices behave better than the digital nets in high dimensional problems.

Both the digital nets and the integration lattices estimators converge at the theoretical (worst case) rate of $O(n^{-1}(\ln n)^t)$. This convergence rate is certainly better than the MC rate $O(n^{-1/2})$ *asymptotically*, but this superiority is practical only for small t . For example, for $t = 10$ already, to have $n^{-1}(\ln n)^t < n^{-1/2}$ one needs $n \approx 10^{39}$. Fortunately, the QMC convergence rate in practice is much better than what is suggested by the theoretical upper bound of $O(n^{-1}(\ln n)^t)$. While this phenomenon is not entirely understood, one explanation comes from the concept of effective dimension of f . This concept is very similar to the effective dimensionality of the correlation matrix as obtained through the principal component analysis. In applications, in particular in financial problems, the effective dimensionality t_{eff} is much less than the nominal dimensionality t , which results in a much faster convergence rate $O(n^{-1}(\ln n)^{t_{\text{eff}}}) \ll O(n^{-1}(\ln n)^t)$.

A.2 Korobov Lattice Rule

The integration lattice in the real space \mathbb{R}^t is a set of vectors

$$L_t = \left\{ \mathbf{v} = \sum_{j=1}^t z_j \mathbf{v}_j, \text{ each } z_j \in \mathbf{Z} \right\} \quad (\text{A.2.1})$$

where $\mathbf{v}_1, \dots, \mathbf{v}_t$ are linearly independent vectors in \square^t which form a basis of the lattice, and z_j are integer. A *lattice rule* (of integration) is a rule of the form (A.1.2) and for which the node set $P_n = \{\mathbf{u}_0, \dots, \mathbf{u}_{n-1}\}$ is the intersection of an integration lattice with the unit hypercube: $P_n = L_t \cap [0,1]^t$. In general, a node set in \square^t requires t basis vectors, but some lattices can be generated by a reduced set of r basis vectors, with $r < t$. One can always write

$$P_n = \{((j_1/n_1)\mathbf{v}_1 + \dots + (j_r/n_r)\mathbf{v}_r) \bmod 1 : 0 \leq j_i < n_i \text{ for } i=1, \dots, r\}, \quad (\text{A.2.2})$$

where the reduction modulo 1 is performed coordinate-wise, the \mathbf{v}_i 's are linearly independent *generating vectors*, and $n = n_1 \cdots n_r$. The smallest r for which this holds is called the *rank* of the lattice rule. For reasons to be explained later, we restrict our attention to $r = 1$. For a rule of rank 1, we have

$$P_n = \{(j/n)\mathbf{v} \bmod 1 : 0 \leq j < n\} \quad (\text{A.2.3})$$

for some vector \mathbf{v} . A simple but important special case is a lattice rule of rank 1 with $\mathbf{v} = (1, a, \dots, a^{t-1})$, which is a Korobov rule (Korobov, 1959).

Using a lattice does not guarantee that the points are well-distributed in the unit hypercube. As the very least, we should require that the projection of the lattice L_t over any d -dimensional subspace of \square^t determined by a subset of coordinates $\{i_1, \dots, i_d\} \subseteq \{1, \dots, t\}$ have the same density as L_t , i.e., the corresponding projection of P_n must have n distinct points—as many as P_n itself. A Korobov rule satisfies this requirement if $\gcd(a, n) = 1$, e.g., if n is prime and $1 \leq a < n$, or if n is a power of 2 and a is odd. Even with this condition in place, both n and a have to be chosen carefully in order to achieve good uniformity in P_n . To minimize the integration error, that is, to ensure that Q_n in (A.1.2) is a good estimator of the integral (A.1.1), the point set P_n must display good uniformity not only in \square^t but in any subspace of \square^t . The importance of the last requirement becomes clear if one recalls that in practice the integral (A.1.1) is determined by a few important dimensions that constitute the effective dimensionality of the problem. The search for the optimal values of n and a is computationally intensive; some of the optimal values are provided in

TABLE 10.

TABLE 10
Best a 's for Certain Values of n

N	A
1,021	76
2,039	1,487
4,093	1,516
8,191	5,130
16,381	4,026
32,749	14,251
65,521	8,950
131,071	28,823

Computer experiments show that rules of higher rank do not exhibit sufficiently uniform lower-dimensional projections, and therefore provide no advantage over the best rank-1 rules. In applications, it is unlikely that more complicated rules can do better than the simple Korobov rules with the values of n and a selected from

TABLE 10.

A.3 Implementation in ERM

For any rule of rank 1, P_n can be constructed in a straightforward way by starting with $\mathbf{u} = \mathbf{0}$ and performing $n-1$ iterations of the form $\mathbf{u} = (\mathbf{u} + \mathbf{v}) \bmod 1$, as the formula (A.2.3) suggests. This requires $O(tn)$ additions modulo 1. If the rule is in Korobov form, P_n is in fact equal to the set of all vectors of t successive output values produced by the linear congruential generator (LCG) defined by the recurrence

$$x_i = (ax_{i-1}) \bmod n, \quad z_i = x_i/n, \quad (\text{A.3.1})$$

from all possible initial states x_0 . That is, $\mathbf{u}_i = (z_0, \dots, z_{t-1}) : x_0 = i, 0 \leq i < n$. By itself, the formula (A.3.1) appears to offer no advantage over (A.2.3), since it requires $O(tn)$ multiplications modulo 1 and offers no random access to any particular dimension (for any given n , the components of u_i are calculated recursively). QMC implemented in this manner would result in an inefficient, slow calculation.

In fact, for some particular values of n and a Korobov rules can be implemented in a much more efficient manner. If n is a prime number and a is a *primitive element* modulo n (i.e., $\nu = n-1$ is the smallest positive ν for which $a^\nu \bmod n = 1$), then the corresponding LCG has the maximal period length of $n-1$, and the point set P_n can be constructed as follows: Start with $x_1 = 1$ and generate the sequence $z_1, z_2, \dots, z_{n+t-2}$ via (A.3.1). Along the way, enumerate $\mathbf{u}_1, \dots, \mathbf{u}_{n-1}$, the overlapping vectors of successive values, so that $\mathbf{u}_i = (z_i, z_{i+1}, \dots, z_{i+t-1})$. Then add the vector $\mathbf{u}_0 = \mathbf{0}$. This requires $O(n+t)$ multiplications by a , modulo n , plus some overhead to shift the vector components at each iteration, instead of $O(tn)$ additions in (A.2.3) or $O(tn)$ multiplications in (A.3.1). The values of n and a contained in

TABLE 10 allow this efficient construction in ERM.

Appendix B: Correlation Matrices, Principal Component Analysis and Random Matrix Theory

An important aspect of risk management is the estimation of the correlations between the price movements of different assets and liabilities. The probability of large losses for a portfolio is often dominated by correlated moves of its different constituents. The study of correlation (or covariance) matrices has a long history in finance, and is one of the cornerstones of Markowitz's theory of optimal portfolios.

B.1 Correlation Matrices and Principal Component Analysis

In order to measure these correlations, one often defines the correlation matrix \mathbf{C} between the M instruments in the portfolio. We will assume, without loss of generality, that the instrument returns $\rho_i = \delta x_i / x_i$ have zero means and are rescaled with unit volatility. The matrix elements C_{ij} of \mathbf{C} are obtained as

$$C_{ij} = \langle \rho_i \rho_j \rangle \quad (\text{B.1.1})$$

Since matrix \mathbf{C} is symmetric, the eigenvalues are real numbers. These numbers also must be positive, otherwise it would be possible to find a certain portfolio of the instruments with a negative variance. Following the exposition in Bouchard and Potters (2003), let \mathbf{v}_a be the normalized eigenvector corresponding to eigenvalue λ_a , with $a = 1, \dots, M$. By definition, $\mathbf{C}\mathbf{v}_a = \lambda_a \mathbf{v}_a$. Let us consider a portfolio such that the dollar amount of instrument i is $v_{a,i}$, the i th component of vector \mathbf{v}_a . The variance of the return of such a portfolio is

$$\sigma_a^2 = \left\langle \left(\sum_{i=1}^M v_{a,i} \rho_i \right)^2 \right\rangle = \sum_{i,j=1}^M v_{a,i} v_{a,j} C_{i,j} \equiv \mathbf{v}_a \cdot \mathbf{C}\mathbf{v}_a = \lambda_a. \quad (\text{B.1.2})$$

We see that λ_a is the variance of the portfolio constructed from the weights $v_{a,i}$. Furthermore, using the fact that different eigenvectors are orthogonal, it is easy to see that the correlation of the returns of two portfolios constructed from two different eigenvectors is zero:

$$\left\langle \left(\sum_{i=1}^M v_{a,i} \rho_i \right) \left(\sum_{j=1}^M v_{b,j} \rho_j \right) \right\rangle = \mathbf{v}_b \cdot \mathbf{C}\mathbf{v}_a = 0 \quad (b \neq a). \quad (\text{B.1.3})$$

Thus, we have obtained a set of uncorrelated random returns e_a , which are the returns of the portfolios constructed from the weights $v_{a,i}$:

$$e_a = \sum_{i=1}^M v_{a,i} \rho_i; \quad \langle e_a^2 \rangle = \lambda_a, \quad \langle e_a e_b \rangle = 0 \quad (a \neq b). \quad (\text{B.1.4})$$

Conversely, we can think of the original returns ρ_i as a linear combination of the uncorrelated factors E_a :

$$\rho_i = \sum_{a=1}^M v_{a,i} e_a. \quad (\text{B.1.5})$$

The last equality holds because the transformation matrix $v_{a,i}$ from the original vectors to the eigenvectors is orthogonal. This decomposition is called principal component analysis (PCA): the correlated fluctuations of a set of random variables are decomposed in terms of the fluctuations of underlying uncorrelated factors. In the case of financial returns, the principal components E_a often have an economic interpretation in terms of financial sectors.

As a practical application, Equation (B.1.5) allows us to simulate normal random variables with any given correlation matrix through a set of uncorrelated variables.

B.2 Empirical Correlation Matrices

The correlation matrix \mathbf{C} is estimated (calibrated) by constructing an empirical correlation matrix $\hat{\mathbf{C}}$ from the historical time series of returns $\rho_i(t)$, where i labels the instrument and t the time, $t = 1, \dots, N$:

$$\hat{C}_{i,j} = \frac{1}{N} \sum_{t=1}^N \rho_i(t) \rho_j(t) \quad (\text{B.2.1})$$

A reliable calibration of a correlation matrix is difficult: if one considers M instruments, the correlation matrix contains $M(M-1)/2$ entries, a problem of quadratic complexity. These entries must be determined from M time series of length N ; if N is not very large compared to M , we should expect that the structure of the matrix is dominated by “measurement noise”. From this point of view, it is illuminating to compare the properties of an empirical correlation matrix $\hat{\mathbf{C}}$ to a purely *random* matrix as one could obtain from a finite series of strictly uncorrelated variables. Deviations from the random matrix case will then indicate the presence of true correlations.

Analysis of such random matrices has its origins in nuclear and condensed matter physics. In the limit of very large matrices, the theory of random matrices (Sengupta and Mitra, 1999) allows one to compute the density of eigenvalues of \hat{C} ,

$$\rho(\lambda) = \frac{Q}{2\pi} \frac{\sqrt{(\lambda_{\max} - \lambda)(\lambda - \lambda_{\min})}}{\lambda}, \quad \lambda \in [\lambda_{\min}, \lambda_{\max}], \quad (\text{B.2.2})$$

$$\lambda_{\min}^{\max} = 1 + 1/Q \pm 2\sqrt{1/Q},$$

where $Q = N/M$.

The most important feature predicted by (B.2.2) is that the eigenvalues of a purely random correlation matrix are always spread over a finite interval $[\lambda_{\min}, \lambda_{\max}]$ around 1; eigenvalues vanish above a certain upper edge λ_{\max} .

In PCA analysis of an empirical correlation matrix calculated from M financial time series of finite duration, one will observe a range of eigenvalues: from very small $\lambda \ll 1$ to large $\lambda \approx M$ (the eigenvalues always add up to M , the size of \hat{C} : $\lambda_1 + \lambda_2 + \dots + \lambda_M = M$). It is clear that the very small eigenvalues are a part of “noise” and therefore the corresponding factors should be dropped from the analysis. The crucial issue that remains unresolved in statistics is how to determine the number of factors to keep—that is, how to distinguish “information” from noise. Although a number of “rules of thumb” have been put forward, none provided a clear quantitative criterion. The result (B.2.2) developed in physical science provides the needed rule: all eigenvalues below λ_{\max} belong to the “noise” band and should be dropped from PCA.

It is often the case that a risk management system estimates an empirical correlation matrix and then employs it as is. As we see, a large part of an empirical correlation matrix must be considered as “noise,” and cannot be trusted for risk management (Laloux et al., 1999).

Appendix C: “Stable” Simulation

It is often said that stochastic simulation is incompatible with scenario analysis, in particular with sensitivity testing. In Feldblum’s words (1995, p. 156): if a simulation model is rerun with a different set of assumptions, the results will change, “but one does not know how much of the change stems from the revised assumption, and how much of the change stems from randomness—from the particular realizations produced by the random number generator in each run.” This appendix describes how ERM overcomes this problem and makes scenario analysis completely compatible with stochastic simulation.

Random number generators produce very long streams of i.i.d. $U(0,1)$ random numbers; the whole stream depends on the numeric value of the seed. A different seed will produce a completely different stream, while the same seed will always produce exactly the same stream (that is why it is more precise to describe the numbers in the stream as “pseudo-random”). If the simulation were to use a single “instance” of the underlying random number generator, while the assumptions differed only in the numeric parameters, the solution to our problem would have been easy: we would simply need to seed the generator with the same seed for both sets of assumptions (seeding the generator is normally unavailable in a spreadsheet, yet trivial in a programming language). In each scenario, the output for both sets would be calculated against the same section of the stream of random numbers; the direction of the change and its average value would then be significant even if the change in each scenario were smaller than the variations between scenarios. Unfortunately, this naive scheme completely breaks down whenever we add or remove any sources of risk (investment positions, lines of business, etc.) or simply change the order they are input into the model. Now, the two sets would use different sections of the random stream in each scenario; as a result, the variation between the scenarios could easily dominate the change from the revised assumption.

The solution is to supply each source of risk (i.e., every random variable) with a random stream of its own, that is, with an independent instance of the random number generator. Of course, we cannot seed all those generators with the same number—the variables would then be perfectly correlated. The seed must uniquely depend on the risk. Naturally, each risk can be described in a unique way: this could be CUSIP identifier for investment positions or, say, some standard description of an insurance business line (like “Workers’ Compensation, Accident Year 2000”). ERM translates these unique textual descriptions into numeric seeds with help of a hash function.

A hash function or hash algorithm is a function for examining the input data (usually text) of arbitrary length and producing a relatively short string of digits, an output hash value. The process of computing such a value is known as hashing. The process of hashing has the property that two different inputs are extremely unlikely to hash to the same hash value. If one runs a piece of information—this article for example—through a hash function, and then changes a single letter and runs the information through the function again, the result would be completely different.

The randomness of each risk's idiosyncratic contribution now comes from its own unique underlying stream of i.i.d. $U(0,1)$ numbers. The correlation between risks comes from a shared set of risk factors (see Sections 3.2.1.2 and 3.2.1.5); these fixed factors are generated through Quasi Monte Carlo simulation in such a way that each factor always uses its unique dimension in the "lattice rule" (see Appendices A and B). This technique results in a reproducible and stable Monte Carlo simulation: one can rerun the simulation and be certain that the change stems from the revised assumption rather than the intrinsic randomness of a simulation.

Appendix D: Risk Measures and Simulation Methods

In this appendix, we define VaR and explain its calculation using simulation techniques. However, in order to obtain a complete picture of risk, and introduce risk measures in the decision-making process, we need to use additional statistics reflecting the interaction of the different pieces (business units, portfolios) that lead to the total risk of the company, as well as potential changes in risk due to changes in the composition of the company's business. Marginal and Incremental VaR are related risk measures that can provide this information on the interaction of different pieces of a portfolio.

D.1 VaR, Marginal VaR, and IVaR

Value-at-Risk (VaR) is one of the most important and widely used statistics that measure the potential risk of economic losses. VaR answers the question: What is the minimum amount that the company can expect to lose with a certain probability over a given horizon? In mathematical terms, VaR corresponds to a percentile of the distribution of portfolio P&L. For a given time horizon, the 100 α % VaR, denoted $\text{VaR}(\alpha)$, is the size of loss that will be exceeded with probability $(1 - \alpha)$. Suppose that the loss incurred by a portfolio during the specified period is given by the random variable L , having some (unknown) cumulative distribution function (cdf) F , so that $\text{Prob}(L \leq y) = F(y)$. The portfolio's 100 α % VaR equals the α -th population quantile of L , that is, $\text{VaR}(\alpha) = F^{-1}(\alpha)$.

The **Marginal VaR** of a position (business line, etc.) with respect to a portfolio (company) can be thought of as the amount of risk that the position is adding to the portfolio. In other words, Marginal VaR tells us how the VaR of our portfolio would change if we sold (ran off) or added a specific position. Marginal VaR can be formally defined as the difference between the VaR of the total portfolio and the VaR of the portfolio without the position. It can be easily shown that Marginal VaR is an increasing function of the correlation ρ between the position and the portfolio. When the VaR of the position is much smaller than the VaR of the portfolio, Marginal VaR will be positive when $\rho > 0$, and negative when $\rho < 0$.

Marginal VaR can be used to compute the amount of risk added by an entire position to the total risk of the portfolio. It is an appropriate risk measure in the context of run-off and acquisition decisions. However, we are also interested in the potential effect that buying or selling a relatively small portion of a position would have on the overall risk. For example, in the process of rebalancing a portfolio, we often wish to decrease our holdings by a small amount rather than liquidate the entire position. Since

Marginal VaR can only consider the effect of selling the whole position, it would be an inappropriate measure of risk contribution for this example.

Incremental VaR (IVaR) is a statistic that provides information regarding the sensitivity of VaR to changes in the portfolio holdings. If we denote by IVaR_i the Incremental VaR for the i th position in the portfolio, and by θ_i the percentage change in size of that position, we can approximate the change in VaR by

$$\Delta \text{VaR} = \sum_i \theta_i \text{IVaR}_i \quad (\text{D.1.1})$$

An important difference between IVaR and Marginal VaR is that the IVaRs of the positions add up to the total VaR of the portfolio:

$$\sum_i \text{IVaR}_i = \text{VaR} \quad (\text{D.1.2})$$

This additive property of IVaR has important applications in the allocation of risk to different units (sectors, countries), where the goal is to keep the sum of the risks equal to the total risk.

For a practical calculation of IVaR, we need a more rigorous definition. Let w_i be the amount of money invested in i th position. We define the IVaR of position i as

$$\text{IVaR}_i = w_i \frac{\partial \text{VaR}}{\partial w_i}. \quad (\text{D.1.3})$$

To verify the additive property of IVaR we need to note that VaR is a homogeneous function of order one of the total amount invested. This means that if we double the investments on each position, the VaR of the new portfolio will be twice as large. That is,

$$\text{VaR}(tw_1, tw_2, \dots, tw_n) = t \text{VaR}(w_1, w_2, \dots, w_n). \quad (\text{D.1.4})$$

Then, by Euler's homogeneous function theorem we have that

$$\text{VaR} = \sum_i w_i \frac{\partial \text{VaR}}{\partial w_i} \equiv \sum_i \text{IVaR}_i. \quad (\text{D.1.5})$$

D.2 Simulation Methods and L-Estimators

In the scenario-based simulation approach, $\text{VaR}(\alpha)$ is estimated from a sample of simulated losses that is drawn from F . We denote this estimated value by $\hat{\text{VaR}}(\alpha)$. The sample of portfolio losses provides an empirical approximation \hat{F} to the true loss distribution F . Consider a set of n scenarios and suppose, for ease of exposition, that the likelihood of each scenario is $1/n$. We denote the loss incurred in the k th scenario as L_k , and the k th order statistic of the sample as $L_{(k)}$, so that $L_{(1)} \leq L_{(2)} \leq \dots \leq L_{(n)}$ are the losses L_1, L_2, \dots, L_n sorted in the ascending order. A common way of defining the empirical cdf for the portfolio losses is

$$\hat{F}(y) = \begin{cases} 0, & y < L_{(1)} \\ k/n, & L_{(k)} \leq y < L_{(k+1)} \\ 1, & y \geq L_{(n)} \end{cases} \quad (\text{D.2.1})$$

A popular estimator of $\text{VaR}(\alpha)$ is the sample quantile, i.e., the quantile of the empirical cdf:

$$\hat{\text{VaR}}(\alpha) = \hat{F}^{-1}(\alpha) = L_{(k)} \quad \text{for} \quad \frac{k-1}{n} < \alpha \leq \frac{k}{n}. \quad (\text{D.2.2})$$

We might try to compute IVaR of a position as a numerical derivative of VaR using a set of scenarios and shifting the position by a small amount—in accord with formula (D.1.3). While in a “standard” simulation approach the simulation error would usually be too large to permit a stable estimate of IVaR, it is entirely possible within the “stable” framework described in 0. However, simulation approach allows a much more efficient method of computing IVaR. This method is based on the fact that we can write IVaR in terms of a conditional expectation.

Let us assume that we have calculated the 95% VaR using Monte Carlo simulation with 1000 scenarios. We decide to use the estimator (D.2.2), so that the 95% VaR corresponds to the 950th ordered P&L scenario. Note that VaR is the sum of the P&L for each position in the 950th scenario

$$\hat{\text{VaR}}(\alpha) = L_{(k)} = \sum_i l_i^{(k)}, \quad (\text{D.2.3})$$

where $l_i^{(k)}$ is the P&L of the position i in the k th ordered scenario. If we increase our holdings in one of the positions by a small amount while keeping the rest constant, the resulting portfolio P&L will still be the 950th largest scenario and hence will still correspond to the 95% VaR. In other words, changing the weight of one position by a

small amount will not change the order of the scenarios. Therefore, the change in VaR given a small change of size Δw_i in position i is $\Delta \text{VaR} = \Delta l_i$. For any given position, the P&L distribution is proportional to the size of the position, $l_i = w_i \tilde{l}_i$, where \tilde{l}_i is the P&L per one unit of money. Assuming that VaR is realized only in the 950th scenario we can write:

$$\begin{aligned} w_i \frac{\partial \text{VaR}}{\partial w_i} &= w_i \frac{\Delta \text{VaR}}{\Delta w_i} \\ &= w_i \frac{\Delta l_i}{\Delta w_i} = w_i \frac{\Delta w_i \cdot \tilde{l}_i}{\Delta w_i} = l_i \end{aligned} \quad (\text{D.2.4})$$

Thus, starting from the VaR estimator, we arrive at a very simple estimator for IVaR:

$$\text{IVaR}_i(\alpha) = l_i^{(k)} \quad (\text{D.2.5})$$

Formula (D.2.5) suggests that we can interpret IVaR for a position as the position P&L in the scenario corresponding to the portfolio VaR estimate. Since VaR is in general realized in more than one scenario, we need to average over all the scenarios where the value of the portfolio is equal to VaR. While formula (D.2.5) applies to the estimators of the form (D.2.2) in the simulation framework, one can derive the following exact formula for IVaR:

$$\text{IVaR}_i = \mathbf{E} \left[l_i \mid \sum_i l_i = \text{VaR} \right] \quad (\text{D.2.6})$$

In other words, IVaR_i is the expected P&L of position i given that the total P&L of the portfolio is equal to VaR.

While the closely related estimators of VaR and IVaR of the form (D.2.2) and (D.2.5) are rather simple and convenient, the two have a significant drawback. Since these estimators depend on only one scenario, both computed VaR and IVaR (especially the latter) can be sensitive to the choice of portfolio scenario. The reason is intuitively clear: we discard the valuable information contained in the adjacent scenarios. This insight motivates us to search for better estimators in the form of linear combinations of order statistics. Such robust estimators are known as L-estimators. Within Seabury ERM, VaR and IVaR are estimated as

$$\begin{aligned} \text{VaR}(\alpha) &= \sum_{k=1}^n v_{k,n}(\alpha) L_{(k)}, \\ \text{IVaR}_i(\alpha) &= \sum_{k=1}^n v_{k,n}(\alpha) l_i^{(k)}, \end{aligned} \quad (\text{D.2.7})$$

with the binomial weights

$$v_{k,n}(\alpha) = \binom{n-1}{k-1} \alpha^{k-1} (1-\alpha)^{n-k}, \quad \sum_{k=1}^n v_{k,n}(\alpha) = 1. \quad (\text{D.2.8})$$

The L-estimators (D.2.7) make use of all information available in the simulation sample. It can be shown (Borovkov et al., 1996; Mausser, 2001), that these robust estimators provide a much higher rate of convergence compared to (D.2.2) and especially (D.2.5).

D.3 VaR, Importance Sampling, and Catastrophe Losses

VaR metrics depend on behavior of the tails of risk distributions. For simulated distributions, the tail measures always present a problem. Whether we measure VaR at, say, 99% level as a 99% percentile of the simulated distribution or use a more robust L-estimator (see above), the estimate still depends on just about 1% of all scenarios. If we had a thousand simulated scenarios, the 99% VaR would depend on about 10 scenarios with the highest losses; the variance of the estimate would be quite significant. The 99.9% VaR would be determined by just one scenario with the highest losses; the variance of such an estimate would be impossible to calculate. In order to calculate the tail measures without a prohibitive increase in the number of simulated scenarios, we need a “targeted” simulation that produces many more tail scenarios than suggested by their probability. We also need a way to calculate the statistics against such “targeted” simulation. This is accomplished with a technique known as importance sampling (Kalos and Whitlock, 1986).

The idea is to draw scenarios from the tail of the P&L distribution with a higher frequency than from the body; when calculating risk measures, the higher frequency is compensated by the lower weight of such scenarios. The resulting weighted estimates will have narrower intervals than the result of a straightforward simulation. In general, importance sampling can be applied to any integral of the form (A.1.1), which includes the moments of the random distribution (standard distribution is not of this form, but can be easily calculated from the mean and the second moment).

Importance sampling critically depends on our prior knowledge of the P&L distribution. Of course, if this knowledge were perfect, there would be no need to run any simulation. What is important is to be able to predict in advance which simulation scenarios will likely end up in the tail. Sometimes, as in the case of mixing regular and catastrophe insurance losses, this task becomes relatively simple.

For the sake of exposition, let us assume that an insurance company has enough internal data for their homeowners' policies to model losses due to events that occur at least once in 10 years. The company would like to compliment their loss simulation based on the internal data with a catastrophe loss distribution (itself a result of a simulation) provided by a catastrophe modeling vendor. The catastrophe distribution covers only events that occur no more often than once in 10 years. In a straightforward simulation, for each tenth scenario, a loss would be drawn from the catastrophe distribution and added to the regular loss. In 1,000 simulations, only 100 would have a catastrophic loss. Out of those, about 10 would be rare events—once in 100 years. These 10 large losses would dominate the overall losses in their scenarios; as a result, the 99% VaR would depend on just 10 scenarios. To apply importance sampling to this situation, one could generate 1,000 scenarios with regular losses, and then add a randomly drawn catastrophe loss to *each* of them. Now we have a set of regular scenarios and a set of catastrophe scenarios, 1,000 scenarios each, with the empirical cumulative distribution functions \hat{F}_r and \hat{F}_{cat} (see Equation (D.2.1)). The weighted mix of these two functions, $\hat{F} = 0.9\hat{F}_r + 0.1\hat{F}_{cat}$ can serve as an estimator of the overall cdf. It is easy to see that the 99% VaR calculated with this function will use about 100 catastrophic scenarios, a tenfold increase in the amount of information. The risk measures calculated in this way will have a significantly reduced variance compared with their counterparts estimated from a straightforward simulation.