



**EXTREME**  
**RISK**  
**MANAGEMENT**

**REVOLUTIONARY APPROACHES  
TO EVALUATING AND MEASURING RISK**

**CHRISTINA RAY**

# EXTREME RISK MANAGEMENT

REVOLUTIONARY APPROACHES  
TO EVALUATING AND MEASURING RISK

藏书章

CHRISTINA RAY



New York Chicago San Francisco Lisbon London  
Madrid Mexico City Milan New Delhi  
San Juan Seoul Singapore Sydney Toronto

Copyright © 2010 by The McGraw-Hill Companies, Inc. All rights reserved. Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher.

1 2 3 4 5 6 7 8 9 0 DOC/DOC 1 5 4 3 2 1 0

ISBN 978-0-07-170059-7

MHID 0-07-170059-5

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold with the understanding that neither the author nor the publisher is engaged in rendering legal, accounting, securities trading, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional person should be sought.

—*From a Declaration of Principles Jointly Adopted by a Committee of the American Bar Association and a Committee of Publishers and Associations*

McGraw-Hill books are available at special quantity discounts to use as premiums and sales promotions or for use in corporate training programs. To contact a representative, please e-mail us at [bulksales@mcgraw-hill.com](mailto:bulksales@mcgraw-hill.com).

This book is printed on acid-free paper.

## P R E F A C E

Shortly after September 11, 2001, I was struck by the stories in the press alleging insider trading in the stocks of the airlines involved in the attack. Dr. William Hery, research professor at the Polytechnic Institute of NYU, and I posited the ability to reverse-engineer actionable market intelligence, or MARKINT, that might be useful in counterterrorism from publicly available market information and prices. One of the earliest supporters of MARKINT was Randolph Tauss, director of corporate strategy at Omnis, Inc., who, while a senior government program manager, enthusiastically adopted the concept.

Since then, my scope and interests have widened. Over the last eight years, I have come to see the intelligence community and the financial community as having essentially similar issues with respect to the measurement and mitigation of extreme risk.

However, I've also seen that the members of each community are largely unaware of the applicability of the other's work to their own problems. The financial community is generally unaware of the mathematical sophistication that has been operational in the defense community for a number of years in the form of network-centric operations within the Department of Defense. Similarly, in spite of new emphasis on open-source intelligence, or OSINT (as opposed to intelligence gathered by sensors or clandestine operations), the national security community is largely unaware of the wealth of information and analytical models available in the financial community that might be retasked to its purposes.

Fortunately, this is changing. The financial crisis of 2007–2008 has been a driver of analytical change, and the missions of the financial and intelligence communities have never before been so aligned. Recently,

Dennis Blair, the new director of national intelligence (DNI), said that the “primary near-term security concern of the United States is the global economic crisis and its geopolitical implications.” The President’s Daily Brief (the PDB) has been joined by the daily Economic Intelligence Brief (the “Butterfly Brief”), both of which are classified intelligence community products. The CIA is now publicly advertising employment opportunities for ex-Wall Streeters.

At the same time, failures in risk management protocols and models by top-tier institutions such as Lehman Brothers and Bear Stearns have spurred a search for alternatives. And recognition of the causal chains that led to the crisis has led some theorists to move from backward-looking statistical models to forward-looking causal inference models.

Fortunately, the financial community is just now starting to take a systems view of the elements and interconnections of the causal network that is the global markets in an effort to explain and anticipate “black swan” events (i.e., undirected, unpredicted, and rare events) and “tipping points” (i.e., points at which a previously rare phenomenon becomes dramatically more common), popularized by Nassim Taleb and Malcolm Gladwell, respectively.

This book contains much of what I’ve learned by having a foot in each world. It is meant to be a modern *Art of War* for those who are involved in the pseudo-warfare of trading in the global capital markets. Appropriately, I’ve been immeasurably assisted by experts from each world.

From the financial community, I would like to thank Dr. Robert Mark, the founder of Black Diamond Risk Enterprises and one of the premier experts in the world on risk management and corporate governance. Bob generously suggested some additions on the subject of systemic risk and current best practice in risk management and contributed examples of typical risk reports.

And from the government community, I would also like to thank Dr. Paul Edward Lehner, consulting scientist at the Center for Integrated Intelligence Systems of The MITRE Corporation. Paul provided me with extraordinarily insightful comments (and an education in cutting-edge developments in mathematical psychology) as well as a unique perspective of the complementary analytical techniques used by the private and public sectors.

I’m also very appreciative to Bryan Ware (CEO), Dr. David Daniels (chief scientist), and Linwood Hudson (VP product development) of Digital Sandbox, Inc., of McLean, Virginia, who provided me with examples of risk inference networks used for purposes of national security.

My thanks also to Professor James Moffat, who allowed me to reproduce his fascinating summary of complexity concepts as they apply to information-age military forces, and to Dr. Julien Diard of the Laboratoire de Psychologie et NeuroCognition, Université Pierre Mendès France, who allowed me to do the same with his comprehensive hierarchy of probabilistic modeling formalisms.

And last but not least, my appreciation to Quantum 4D, Inc., and Palisades Corporation, who generously provided me with network analysis and visualization tools with which to create working examples of concepts from this book.

# C O N T E N T S

**Preface**   vii

Chapter 1

**Plausibility versus Probability: Two Worldviews**   1

Chapter 2

**The Evolution of Modern Analytics**   8

Chapter 3

**Natural Selection on Wall Street**   23

Chapter 4

**A Review: Risk Management Metrics, Models, and Best Practice**   32

Chapter 5

**Systemic Risk and Systems Thinking**   61

Chapter 6

**Belief Systems and the Inadvertent Conspiracy: The Global Capital Markets as System**   81

Chapter 7

**Analytic Tradecraft and Model Risk**   94

Chapter 8

**The Future as Forecast: Probability, Likelihood, and Uncertainty**   117

Chapter 9

**An Alternative Path to Actionable Intelligence 144**

Chapter 10

**Solutions: Moving toward a Connectivist Approach 153**

Chapter 11

**An Introduction to Causality: Theory, Models,  
and Inference 165**

Chapter 12

**Bayesian Inference Networks: Estimating Vulnerability,  
Consequences, and Likelihood 175**

Chapter 13

**Securities Valuation, Risk Measurement,  
and Risk Inference Networks 189**

Chapter 14

**Risk Fusion and Supermodels: A Framework  
for Enterprise Risk Management 206**

Chapter 15

**Inferring Causality from Historical Market Behavior 216**

Chapter 16

**Sensemaking for Warnings: Reverse-Engineering  
Market Intelligence 230**

Chapter 17

**The United States as Enterprise: Implications for  
National Policy and Security 245**

**Notes 263**

**Index 279**



# Plausibility versus Probability: Two Worldviews

Over the last three decades or so, sophisticated financial modeling has been almost exclusively statistical in nature. The ready availability of massive amounts of historical market data has fueled the creation of valuation and risk measurement models built on concepts such as association, correlation, and likelihood.

All these models create implicit forecasts, that is, estimates of expected and possible future scenarios for a security or a portfolio of securities. Most often, these forecasts are based on the assumption that the future market behavior is well represented by the past.

However, this stochastic approach implies a worldview that ignores *causality* in favor of *correlation*. In this world, it doesn't matter whether gold prices increased because interest rates decreased or vice versa. It also doesn't matter whether the prices of a utility stock and an airline stock are directly related in some fashion or whether, instead, they are both driven by a common dependence on fuel prices. This world is a supremely efficient world as well: all prices reflect new information immediately, and that information is transmitted instantaneously around the globe.

However, intuition belies these notions. Traders and portfolio managers know that events drive prices. Catalysts such as the release of an economic indicator or an earnings report drive prices, and chain reactions precipitated by an important event can take a finite amount of time to propagate.

Such statistical models were often sufficient in the past, when the volume and complexity of derivative instruments were far lower than they are today. But now, the value and risk of popular instruments such as options on credit derivatives and complex asset-backed securities increasingly depend on the modeling of low-probability, high-consequence events. If the models used are not adequate for the task of anticipating such high-consequence events, massive losses and market disruptions can

occur. Certainly the financial disruptions that began in 2007–2008 are abundant evidence of such failures.

But as the old saying goes, “Correlation is not causation.” The alternative to a statistical model is a causal model that explicitly creates an alternative worldview, one in which cause and effect are modeled in logical or temporal order.

This alternative world is one in which *plausibility* rather than *probability* is modeled. The consequences and likelihood of events that have never before occurred but that can be reasonably anticipated (as a consequence of other events) are included in the quantitative models. Such modeling is the forte of the intelligence community and those responsible for national security, who must create metrics and construct solutions for threats that have never before occurred.

Plausibility can be determined from a mixture of expert opinion, hard facts, and historical experience. Although the structure of any causal model may be guided by the insights of human experts, it need not be strictly an expert system. Instead, through a process of causal inference, past history can be used to validate and inform the model. A causal model is not necessarily deterministic; it can allow for uncertainty. Ideally, causal inference facilitates the integration of substantive knowledge with statistical data to refine the former and interpret the latter.

Such causal models are used in other disciplines, most notably epidemiology and decision science. They are little used in finance, with the notable exception of the measurement of operational risk (i.e., the risk of loss due to human error). Causal models are nearly absent from<sup>1</sup> quantitative modeling for purposes of instrument valuation or market and credit risk measurement.

The preference that quantitative analysts have for “frequentist” or probabilistic models over causal models over the last three decades is understandable for a number of reasons.

First, such models are relatively easy to create and implement, using financial theories (such as modern portfolio theory) that are already well accepted and in the public domain.

Also, until recently, neither the mathematical language nor the technical tools that might facilitate the creation of causal models existed. Although the financial community commenced serious quantitative modeling in the 1970s, it wasn’t until the mid-1980s that much substantive work was done on causal models, even within the academic community.

Thus, the creation of rigorous theory, methodologies, and a language of causality that might have facilitated such model building did not exist at the time the financial community was choosing its path. Perhaps more

important, even if such models had been created, the data required to inform them were usually insufficiently granular, synchronized, and properly organized for use in a causal inference process.

However, now, in the words of Judea Pearl, a leader in this field, “Put simply, causality has been mathematized.” At the same time, certain technological innovations have made causal inference practical in the financial arena.

Consider one of the key questions in causal inference: How can one distinguish between mere correlation and cause and effect? When the sun rises and the cock crows, was one of these two events the catalyst for the other, or were they both the consequence of a third event?

One of the best methods of validating causal relationships is via experimentation. We can wake up the cock at 3 a.m. and see if this causes the sun to rise. Or an experiment can be designed to eliminate all variables but one: for example, in medical trials, the effect of the drug on a patient. To produce valid results, such an experiment would probably contain key features used in causal modeling, such as randomization (e.g., patients are randomly selected to receive an experimental drug or a placebo) and elimination of exogenous factors (e.g., variations in age or sex).

Fortunately, in finance, the capital markets are a laboratory that continuously provides us with natural experiments. Thus, rather than using historical market prices in statistical analyses, we can use them in causal inference models. Every day, traders receive information about catalytic events that move markets and are able to observe the synchronous or subsequent effects of those catalysts.

Technological advances now make the observation of these natural experiments both possible and practical. Formerly, end-of-day data were relatively useless for determining causation because so many important events occur during the course of a trading day. Just as in a medical trial, when there are multiple variables, reliable causal inferences are exceedingly difficult to make.

Only in the last few years has commercial software become available that is capable of capturing event data and synchronizing those data with real-time market data of the highest granularity. This synchronized information gives us the means to learn from one controlled experiment at a time, even if the experiments last just seconds.

Although many events occur in the course of a trading day, few of them occur simultaneously, where *simultaneous* is defined as occurring within the same very small window of time. For example, we might capture the earliest moment at which an earnings report became public or a report on crude oil inventory was released. If we then examine the real-time

behavior of stock or oil prices in the seconds to minutes after the release, we can form opinions about how such an event drives prices.

Besides potentially providing better estimates of value and risk, causal models may be more intuitive and understandable by risk managers and portfolio managers than statistical models are. For example, the language of causality is a natural language for risk management. Examples of causal concepts are influence, ignorability, disturbance, effect, confounding, intervention, and explanation.

The graphic representations that substitute for mathematical equations lend themselves well to financial applications. As Pearl points out, there is no analog in algebra or statistics to the causal operator “given that I do,” that is, the effect of a deliberate action on the outcome of the analysis. However, these representations lend themselves well to programming. Computer code *does* allow such operators; the statement  $A = B$  is a substitution rather than a statement about an inviolate relationship between A and B.

Similarly, hedge positions can be considered “interventions” that can block certain paths: those that lead to undesirable outcomes, such as very large losses. Such a hedge might be a security that is already in a portfolio or, alternatively, an exogenous variable that drives changes in one or more securities in the portfolio.

Further, the identification of hierarchically organized causes lends itself very naturally to the identification of systematic and specific risks as required by the Basel II accord. Such methods may provide results that are far superior to those provided by statistical methodologies such as principal component analysis,<sup>2</sup> the results of which can be degraded by spurious correlations without expert intervention.

Causal models also provide a natural framework for the estimation of two key risk measures for which no industry-standard methods yet exist: economic capital (the amount of capital required to ensure the continued existence of the enterprise to a very high degree of confidence) and enterprise risk (the risk to an enterprise from all sources of risk). In such models, expert opinion can be integrated with historical behavior to systematically generate all plausible future scenarios, estimate their likelihood, and measure their consequences.

All else being equal, a causal approach is preferable to a statistical approach for several reasons.

First, a causal approach allows a more general solution. A statistical solution can be simulated, albeit inefficiently, using a causal network that includes an error component. However, the reverse is not true.

Second, causal networks do not require extensive historical data for all the securities and instruments in a portfolio. Causal models can be used

even when history is not a reliable indicator of the future—for example, when a shift in risk regime has occurred or when new risk factors such as changing regulatory policy are expected to have a significant impact.

Causal networks can also be allowed to have a specific order in which events occur or a temporal component suitable for high-frequency trading and real-time risk management. Forecasts of consequential behavior produced in sufficient time to execute a trade can be used in automated, algorithmic trading. Further, observed market behavior that is time dependent (e.g., volatility clusters and jump diffusion processes) might be more easily explained in terms of causal models than it is by statistical models. Also, instead of relying on solutions such as GARCH<sup>3</sup> methods or stochastic volatility models<sup>4</sup> to calibrate observations to history, such observations might be explained in terms of the observable, noninstantaneous effects of traders' and portfolio managers' behavior.<sup>5</sup>

A causal approach can use *all* available information to inform the model, not just historical pricing data. In the terminology of the intelligence community, this is the use of “all source intelligence.” For example, additional fundamental information might be used to inform (or override) certain causal relationships. The sensitivity of an airline to the price of fuel might be independently modeled by a fundamental equity analyst and then compared to the relationship inferred by the causal model. Or if a publicly traded home builder has never before hedged its interest-rate exposure but has just started such a program, the past dependence of the company's stock price on interest rates might be overridden.

A causal approach is far more dynamic than a statistical approach because it allows the introduction of *prior knowledge*. A forecast of one-day risk is substantially different one second *after* the release of the monthly unemployment statistics from what it was one second *before* that release, based on knowledge of both market expectations and the actual news. In the language of causal modeling, these are the *prior* and *posterior distributions*.

Most important of all, a causal structure provides far more transparency than do statistical parameters. The graphical language of causal modeling reveals the fundamental relationships assumed by experts and inferred from data and lends itself to the use of visualization tools that enhance clarity and aid human cognition.

The process of building such a model also removes some of the intellectual barriers between the front office and the middle office and between technical analysis and fundamental analysis. Causal relationships that can be understood and vetted by human experts with multiple areas of expertise are far more likely to be repeated in the future.

Clearly, causal models are somewhat more difficult to implement and to inform than are statistical models. However, when they are used for

certain purposes, such as valuing complex derivative instruments, estimating extreme or real-time portfolio risk, or designing an optimal hedging strategy, they are well worth the effort.

For example, one of their major advantages is the ability to perform discrete-time and discrete-outcome modeling. Although common statistical methodologies such as copula approaches are mathematically elegant, they often implicitly eliminate the granularity, asymmetry, and noncontinuous behavior that are interesting features (and opportunities for profit) of real markets. By doing so, they may substantially over- or underestimate value or risk, particularly for instruments with a narrow payoff window, such as  $n$ th-to-default tranches in collateralized debt obligations,<sup>6</sup> or in strategies such as calendar or price spreads in options.

What a causal approach lacks in computational elegance it may make up for in accuracy. Consider a situation in which Treasury bond traders are split 50/50 on whether the Treasury will announce an auction of 30-year bonds. This is a binary event: it will occur, or it will not; the yield curve will flatten or steepen. A realistic forecast of changes in 30-year bond yields just after the announcement is likely to be bimodal because there is no neutral event.

The benefit of a causal model is its ability to generate many plausible scenarios in a systematic fashion. The likelihood of some of these scenarios may be higher than in a random-walk world; that of others may be lower. Markets may have “hot spots” and “cold spots”: scenarios in which a convergence of certain chain reactions is likely to have major market repercussions or, conversely, scenarios that are virtually impossible.

This alternative forecast of the future, in which the distributions of possible outcomes can be granular, be asymmetrical, and have extreme outcomes, has profound implications for financial engineering, portfolio management, risk management, and even decision science. Clearly, a set of possible future scenarios substantially different from those created using continuous, normally distributed variables, suggests radically different results for all kinds of estimates.

Specifically, valuation models, particularly those for securities or complex derivative instruments, will produce results that are substantially different from the results of standard models that assume normality, symmetry, and outcomes that are in line with historical experience. Portfolio optimization and performance attribution models are similarly affected. The interactions between the securities in the portfolio may be poorly described by statistical measures such as correlation, and an ideal portfolio (i.e., one with an optimum risk-return profile) constructed using such

scenarios might look quite different from one constructed using more traditional methods.

Most important, risk measurement models based on causal methods may estimate risk to have a magnitude that is either far greater or far less—as well as less continuous—than that estimated using traditional stochastic models. Certain outcomes that were formerly assumed to be virtually impossible must now be considered, whereas others are now less likely than before. At the heart of all risk measurement (and in fact all financial engineering) is the ability to generate all plausible alternative scenarios and estimate their likelihood. The sensitivity of a portfolio, an enterprise, or even the global capital markets themselves to the most extreme of these scenarios provides a systematic method for generating *stress tests* (i.e., measures of the consequences of a particular scenario) and ensuring the continued existence of the system as we know it.

The use of causal methods also provides solutions; they can be used to mitigate as well as measure risk. They provide a method for inserting circuit breakers into a portfolio or a banking system to subvert the most catastrophic outcomes. For example, a portfolio manager might purchase far out-of-the-money calls on crude oil to hedge the risk of large declines in the price of airline and hospitality stocks, or a regulator might modify capital requirements or position-limit rules.

The ultimate goal of enterprise risk management is as a quantitative decision-making tool. Ultimately, the use of causal methods facilitates the highest-level goal of risk management: decision making by senior management. An understanding of the possible future paths that might trigger tipping points and lead to catastrophic outcomes can assist C-suite executives in optimizing their business strategies on a risk-adjusted basis.

## SUMMARY

In this chapter, we contrasted frequentist and causal approaches to risk management and their utility in measuring and mitigating extreme risk and optimizing decision making.

## WHAT'S NEXT

In the next chapter, we will relate key milestones in the evolution of risk management philosophy and describe the most recent and revolutionary innovations in quantitative decision making.

## The Evolution of Modern Analytics

The recent journey toward the current state of quantitative finance owes as much to philosophy as it does to mathematics. It has been said that a human can't visualize a number of objects greater than five without breaking them into smaller groups, such as two sets of three or three sets of two. For more complicated problems, humans have always had to create an abstraction or model of how things work.

As eloquently described by Peter Bernstein in *Against the Gods: The Remarkable Story of Risk*<sup>1</sup> scientists of the seventeenth and eighteenth centuries considered the problem of decision making under uncertainty. Some of the greatest mathematicians in history (including Isaac Newton) were tasked by their patrons with solving gaming problems, with profit as the motive.

Such problems provided the perfect thought experiment, involving as they did both chance and preferences. In a sense, they forced the evolution of the financial markets, since those who used these insights to inform their decisions prospered, whereas those who did not became extinct.

This process of financial natural selection continues today. It's a short leap from the gaming table to the trading room. Broker-dealers and hedge funds that were able to adequately manage risk survived and prospered during the financial crisis of 2008, whereas those that were not (including Bear Stearns, Lehman Brothers, and AIG) did not.

Along the evolutionary path are two historical milestones that were separated by more than 200 years but are suddenly receiving more attention from the mainstream of economic thought nearly simultaneously.

### **DANIEL BERNOULLI AND A NEW THEORY ON THE MEASUREMENT OF RISK**

In 1954, spurred by a spate of recent references to this work, the journal *Econometrica*<sup>2</sup> published the first English translation from the original



Latin of Daniel Bernoulli's 1738 *Specimen theoriae novae de mensura sortis*, or *Exposition of a New Theory on the Measurement of Risk*.<sup>3</sup>

Bernoulli is far more famous as a physicist for his contributions to fluid mechanics, including the basic principles that allow aircraft to fly. *Exposition* was considered an exceedingly minor scientific contribution until its rediscovery by economists, evolutionary biologists, computer scientists, and others in the second half of the twentieth century. Although Dr. Louise Sommer, the translator, explicitly attempted to retain the article's eighteenth-century flavor, the concepts and language of the article are astoundingly up to date.

This amazingly concise and complete piece of work not only discussed basic probability theory but also put forth the concepts of individual preferences and utility functions, which together form the foundation of a framework for cutting-edge topics such as quantitative decision support<sup>4</sup> and enterprise risk management (ERM).

For example, in discussing optimum wagering decisions, Bernoulli rejected the notion that the expected outcome (based on probabilities) should always be maximized and introduced sophisticated concepts such as risk premia<sup>5</sup> and utility functions:<sup>6</sup>

Somehow a very poor fellow obtains a lottery ticket that will yield with equal probability either nothing or twenty thousand ducats. Will this man evaluate his chance of winning at ten thousand ducats? Would he not be ill-advised to sell this lottery ticket for nine thousand ducats? To me it seems that the answer is in the negative. On the other hand, I am inclined to believe that a rich man would be ill-advised to refuse to buy the lottery ticket for nine thousand ducats. If I am not wrong then it seems clear that all men cannot use the same rule to evaluate the gamble.

Bernoulli went on to more formally introduce key concepts such as quantitative measurement of risk, diversification, hedging, and, perhaps most important, the concept of maximizing utility rather than expected return:

Thus it becomes evident that no valid measurement of the value of a *risk* can be obtained without consideration being given to its *utility*, that is to say, the utility of whatever gain accrues to the individual or, conversely, how much profit is required to yield a given utility . . . the determination of the value of an item must not be based upon its price, but rather the utility it yields.

Long before the advent of behavioral economics, Bernoulli proposed the famous St. Petersburg Paradox, first suggested to him by his distinguished cousin Nicholas Bernoulli. In this simple but elegant thought